1, TITLE: Unsupervised Learning of Probably Symmetric Deformable 3D Objects From Images in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Unsupervised_Learning_of_Probably_Symmetric_Deformable_3D_Objects_From_Images_CVPR_2020_paper.html
AUTHORS:	Shangzhe Wu, Christian Rupprecht, Andrea Vedaldi
HIGHLIGHT:	We propose a method to learn 3D deformable object categories from raw single-view images, without external supervision.

2, TITLE: Footprints and Free Space From a Single Color Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Watson_Footprints_and_Free_Space_From_a_Single_Color_Image_CVPR_2020_paper.html
AUTHORS:	Jamie Watson, Michael Firman, Aron Monszpart, Gabriel J. Brostow
HIGHLIGHT:	We introduce a model to predict the geometry of both visible and occluded traversable surfaces, given a single RGB image as input.

3, TITLE: Dynamic Fluid Surface Reconstruction Using Deep Neural Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Thapa_Dynamic_Fluid_Surface_Reconstruction_Using_Deep_Neural_Network_CVPR_2020_paper.html
AUTHORS:	Simron Thapa, Nianyi Li, Jinwei Ye
HIGHLIGHT:	Here we present a learning-based single-image approach for 3D fluid surface reconstruction.

4, TITLE: CvxNet: Learnable Convex Decomposition
http://openaccess.thecvf.com/content_CVPR_2020/html/Deng_CvxNet_Learnable_Convex_Decomposition_CVPR_2020_paper.html
AUTHORS:	Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, Andrea Tagliasacchi
HIGHLIGHT:	We introduce a network architecture to represent a low dimensional family of convexes.

5, TITLE: BSP-Net: Generating Compact Meshes via Binary Space Partitioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_BSP-Net_Generating_Compact_Meshes_via_Binary_Space_Partitioning_CVPR_2020_paper.html
AUTHORS:	Zhiqin Chen, Andrea Tagliasacchi, Hao Zhang
HIGHLIGHT:	The core ingredient of BSP is an operation for recursive subdivision of space to obtain convex sets. By exploiting this property, we devise BSP-Net, a network that learns to represent a 3D shape via convex decomposition.

6, TITLE: Total3DUnderstanding: Joint Layout, Object Pose and Mesh Reconstruction for Indoor Scenes From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Nie_Total3DUnderstanding_Joint_Layout_Object_Pose_and_Mesh_Reconstruction_for_Indoor_CVPR_2020_paper.html
AUTHORS:	Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, Jian Jun Zhang
HIGHLIGHT:	In this paper, we bridge the gap between understanding and reconstruction, and propose an end-to-end solution to jointly reconstruct room layout, object bounding boxes and meshes from a single image.

7, TITLE: Generating and Exploiting Probabilistic Monocular Depth Estimates
http://openaccess.thecvf.com/content_CVPR_2020/html/Xia_Generating_and_Exploiting_Probabilistic_Monocular_Depth_Estimates_CVPR_2020_paper.html
AUTHORS:	Zhihao Xia, Patrick Sullivan, Ayan Chakrabarti
HIGHLIGHT:	Instead, we propose a versatile task-agnostic monocular model that outputs a probability distribution over scene depth given an input color image, as a sample approximation of outputs from a patch-wise conditional VAE.

8, TITLE: Neural Cages for Detail-Preserving 3D Deformations
http://openaccess.thecvf.com/content_CVPR_2020/html/Yifan_Neural_Cages_for_Detail-Preserving_3D_Deformations_CVPR_2020_paper.html
AUTHORS:	Wang Yifan, Noam Aigerman, Vladimir G. Kim, Siddhartha Chaudhuri, Olga Sorkine-Hornung
HIGHLIGHT:	We propose a novel learnable representation for detail preserving shape deformation.

9, TITLE: PIFuHD: Multi-Level Pixel-Aligned Implicit Function for High-Resolution 3D Human Digitization
http://openaccess.thecvf.com/content_CVPR_2020/html/Saito_PIFuHD_Multi-Level_Pixel-Aligned_Implicit_Function_for_High-Resolution_3D_Human_Digitization_CVPR_2020_paper.html
AUTHORS:	Shunsuke Saito, Tomas Simon, Jason Saragih, Hanbyul Joo
HIGHLIGHT:	Due to memory limitations in current hardware, previous approaches tend to take low resolution images as input to cover large spatial context, and produce less precise (or low resolution) 3D estimates as a result. We address this limitation by formulating a multi-level architecture that is end-to-end trainable.

10, TITLE:	A Lighting-Invariant Point Processor for Shading

1

http://openaccess.thecvf.com/content_CVPR_2020/html/Heal_A_Lighting-
Invariant_Point_Processor_for_Shading_CVPR_2020_paper.html
AUTHORS: Kathryn Heal, Jialiang Wang, Steven J. Gortler, Todd Zickler
HIGHLIGHT: We describe the geometry of this variety, and we introduce a concise feedforward model that computes an explicit, differentiable approximation of the variety from the intensity and its derivatives at any single image point.

11, TITLE: ActiveMoCap: Optimized Viewpoint Selection for Active Human Motion Capture
http://openaccess.thecvf.com/content_CVPR_2020/html/Kiciroglu_ActiveMoCap_Optimized_Viewpoint_Selection_for_Active_Human_Motion_Capture_CVPR_2020_paper.html
AUTHORS: Sena Kiciroglu, Helge Rhodin, Sudipta N. Sinha, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: Specifically, given a short video sequence, we introduce an algorithm that predicts which viewpoints should be chosen to capture future frames so as to maximize 3D human pose estimation accuracy.

12, TITLE: Peek-a-Boo: Occlusion Reasoning in Indoor Scenes With Plane Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Peek-a-
Boo_Occlusion_Reasoning_in_Indoor_Scenes_With_Plane_Representations_CVPR_2020_paper.html
AUTHORS: Ziyu Jiang, Buyu Liu, Samuel Schulter, Zhangyang Wang, Manmohan Chandraker
HIGHLIGHT: We address the challenging task of occlusion-aware indoor 3D scene understanding.

13, TITLE: Multi-Modal Domain Adaptation for Fine-Grained Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Munro_Multi-Modal_Domain_Adaptation_for_Fine-
Grained_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Jonathan Munro, Dima Damen
HIGHLIGHT: In this work we exploit the correspondence of modalities as a self-supervised alignment approach for UDA in addition to adversarial alignment (Fig. 1).

14, TITLE: Evolving Losses for Unsupervised Video Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Piergiovanni_Evolving_Losses_for_Unsupervised_Video_Representation_Learning_CVPR_2020_paper.html
AUTHORS: AJ Piergiovanni, Anelia Angelova, Michael S. Ryoo
HIGHLIGHT: We present a new method to learn video representations from large-scale unlabeled video data.

15, TITLE: Disentangling and Unifying Graph Convolutions for Skeleton-Based Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Disentangling_and_Unifying_Graph_Convolutions_for_Skeleton-
Based_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, Wanli Ouyang
HIGHLIGHT: In this work, we present (1) a simple method to disentangle multi-scale graph convolutions and (2) a unified spatial-temporal graph convolutional operator named G3D.

16, TITLE: A Multigrid Method for Efficiently Training Video Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_A_Multigrid_Method_for_Efficiently_Training_Video_Models_CVPR_2020_paper.html
AUTHORS: Chao-Yuan Wu, Ross Girshick, Kaiming He, Christoph Feichtenhofer, Philipp Krahenbuhl
HIGHLIGHT: Inspired by multigrid methods in numerical optimization, we propose to use variable mini-batch shapes with different spatial-temporal resolutions that are varied according to a schedule.

17, TITLE: Ego-Topo: Environment Affordances From Egocentric Video
http://openaccess.thecvf.com/content_CVPR_2020/html/Nagarajan_Ego-
Topo_Environment_Affordances_From_Egocentric_Video_CVPR_2020_paper.html
AUTHORS: Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, Kristen Grauman
HIGHLIGHT: We introduce a model for environment affordances that is learned directly from egocentric video.

18, TITLE: Generative Hybrid Representations for Activity Forecasting With No-Regret Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Guan_Generative_Hybrid_Representations_for_Activity_Forecasting_With_
No-Regret_Learning_CVPR_2020_paper.html
AUTHORS: Jiaqi Guan, Ye Yuan, Kris M. Kitani, Nicholas Rhinehart
HIGHLIGHT: In this work, we develop an efficient deep generative model to jointly forecast a person's future discrete actions and continuous motions.

19, TITLE: Skeleton-Based Action Recognition With Shift Graph Convolutional Network

http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Skeleton-Based_Action_Recognition_With_Shift_Graph_Convolutional_Network_CVPR_2020_paper.html
AUTHORS: Ke Cheng, Yifan Zhang, Xiangyu He, Weihan Chen, Jian Cheng, Hanqing Lu
HIGHLIGHT: In this paper, we propose a novel shift graph convolutional network (Shift-GCN) to overcome both shortcomings.

20, TITLE: Predicting Goal-Directed Human Attention Using Inverse Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Predicting_Goal-Directed_Human_Attention_Using_Inverse_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS: Zhibo Yang, Lihan Huang, Yupei Chen, Zijun Wei, Seoyoung Ahn, Gregory Zelinsky, Dimitris Samaras, Minh Hoai
HIGHLIGHT: We propose the first inverse reinforcement learning (IRL) model to learn the internal reward function and policy used by humans during visual search.

21, TITLE: X3D: Expanding Architectures for Efficient Video Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Feichtenhofer_X3D_Expanding_Architectures_for_Efficient_Video_Recognition_CVPR_2020_paper.html
AUTHORS: Christoph Feichtenhofer
HIGHLIGHT: This paper presents X3D, a family of efficient video networks that progressively expand a tiny 2D image classification architecture along multiple network axes, in space, time, width and depth.

22, TITLE: Dynamic Multiscale Graph Neural Networks for 3D Skeleton Based Human Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Dynamic_Multiscale_Graph_Neural_Networks_for_3D_Skeleton_Based_Human_CVPR_2020_paper.html
AUTHORS: Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, Qi Tian
HIGHLIGHT: We propose novel dynamic multiscale graph neural networks (DMGNN) to predict 3D skeleton-based human motions.

23, TITLE: Use the Force, Luke! Learning to Predict Physical Forces by Simulating Effects
http://openaccess.thecvf.com/content_CVPR_2020/html/Ehsani_Use_the_Force_Luke_Learning_to_Predict_Physical_Forces_by_CVPR_2020_paper.html
AUTHORS: Kiana Ehsani, Shubham Tulsiani, Saurabh Gupta, Ali Farhadi, Abhinav Gupta
HIGHLIGHT: In this paper, we take a step towards more physical understanding of actions.

24, TITLE: DaST: Data-Free Substitute Training for Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_DaST_Data-Free_Substitute_Training_for_Adversarial_Attacks_CVPR_2020_paper.html
AUTHORS: Mingyi Zhou, Jing Wu, Yipeng Liu, Shuaicheng Liu, Ce Zhu
HIGHLIGHT: In this paper, we propose a data-free substitute training method (DaST) to obtain substitute models for adversarial black-box attacks without the requirement of any real data.

25, TITLE: Towards Verifying Robustness of Neural Networks Against A Family of Semantic Perturbations
http://openaccess.thecvf.com/content_CVPR_2020/html/Mohapatra_Towards_Verifying_Robustness_of_Neural_Networks_Against_A_Family_of_CVPR_2020_paper.html
AUTHORS: Jeet Mohapatra, Tsui-Wei Weng, Pin-Yu Chen, Sijia Liu, Luca Daniel
HIGHLIGHT: To bridge this gap, we propose Semantify-NN, a model-agnostic and generic robustness verification approach against semantic perturbations for neural networks.

26, TITLE: The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_The_Secret_Revealer_Generative_Model-Inversion_Attacks_Against_Deep_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, Dawn Song
HIGHLIGHT: Here we present a novel attack method, termed the generative model-inversion attack, which can invert deep neural networks with high success rates.

27, TITLE: A Self-supervised Approach for Adversarial Robustness
http://openaccess.thecvf.com/content_CVPR_2020/html/Naseer_A_Self-supervised_Approach_for_Adversarial_Robustness_CVPR_2020_paper.html
AUTHORS: Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Fatih Porikli
HIGHLIGHT: In this paper, we take the first step to combine the benefits of both approaches and propose a self-supervised adversarial training mechanism in the input space.

28, TITLE: Adversarial Vertex Mixup: Toward Better Adversarially Robust Generalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Adversarial_Vertex_Mixup_Toward_Better_Adversarially_Robust_Generalization_CVPR_2020_paper.html
AUTHORS: Saehyung Lee, Hyungyu Lee, Sungroh Yoon
HIGHLIGHT: In this paper, we identify Adversarial Feature Overfitting (AFO), which may cause poor adversarially robust generalization, and we show that adversarial training can overshoot the optimal point in terms of robust generalization, leading to AFO in our simple Gaussian model.

29, TITLE: How Does Noise Help Robustness? Explanation and Exploration under the Neural SDE Framework
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_How_Does_Noise_Help_Robustness_Explanation_and_Exploration_under_the_CVPR_2020_paper.html
AUTHORS: Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, Cho-Jui Hsieh
HIGHLIGHT: In this paper, we propose a new continuous neural network framework called Neural Stochastic Differential Equation (Neural SDE), which naturally incorporates various commonly used regularization mechanisms based on random noise injection.

30, TITLE: Unpaired Image Super-Resolution Using Pseudo-Supervision
http://openaccess.thecvf.com/content_CVPR_2020/html/Maeda_Unpaired_Image_Super-Resolution_Using_Pseudo-Supervision_CVPR_2020_paper.html
AUTHORS: Shunta Maeda
HIGHLIGHT: In this paper, we propose an unpaired SR method using a generative adversarial network that does not require a paired/aligned training dataset.

31, TITLE: Universal Litmus Patterns: Revealing Backdoor Attacks in CNNs
http://openaccess.thecvf.com/content_CVPR_2020/html/Kolouri_Universal_Litmus_Patterns_Revealing_Backdoor_Attacks_in_CNNs_CVPR_2020_paper.html
AUTHORS: Soheil Kolouri, Aniruddha Saha, Hamed Pirsiavash, Heiko Hoffmann
HIGHLIGHT: In this paper, we introduce a benchmark technique for detecting backdoor attacks (aka Trojan attacks) on deep convolutional neural networks (CNNs).

32, TITLE: Robustness Guarantees for Deep Neural Networks on Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Robustness_Guarantees_for_Deep_Neural_Networks_on_Videos_CVPR_2020_paper.html
AUTHORS: Min Wu, Marta Kwiatkowska
HIGHLIGHT: In this paper, we consider the robustness of deep neural networks on videos, which comprise both the spatial features of individual frames extracted by a convolutional neural network and the temporal dynamics between adjacent frames captured by a recurrent neural network.

33, TITLE: Benchmarking Adversarial Robustness on Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Benchmarking_Adversarial_Robustness_on_Image_Classification_CVPR_2020_paper.html
AUTHORS: Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, Jun Zhu
HIGHLIGHT: In this paper, we establish a comprehensive, rigorous, and coherent benchmark to evaluate adversarial robustness on image classification tasks.

34, TITLE: What It Thinks Is Important Is Important: Robustness Transfers Through Input Gradients
http://openaccess.thecvf.com/content_CVPR_2020/html/Chan_What_It_Thinks_Is_Important_Is_Important_Robustness_Transfers_Through_CVPR_2020_paper.html
AUTHORS: Alvin Chan, Yi Tay, Yew-Soon Ong
HIGHLIGHT: Using only natural images, we show here that training a student model's input gradients to match those of a robust teacher model can gain robustness close to a strong baseline that is robustly trained from scratch.

35, TITLE: Transferable, Controllable, and Inconspicuous Adversarial Attacks on Person Re-identification With Deep Mis-Ranking
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Transferable_Controllable_and_Inconspicuous_Adversarial_Attacks_on_Person_Re-identification_With_CVPR_2020_paper.html
AUTHORS: Hongjun Wang, Guangrun Wang, Ya Li, Dongyu Zhang, Liang Lin
HIGHLIGHT: In this work, we examine the insecurity of current best-performing ReID models by proposing a learning-to-mis-rank formulation to perturb the ranking of the system output.

36, TITLE: Video Modeling With Correlation Networks

http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Video_Modeling_With_Correlation_Networks_CVPR_2020_paper.html
AUTHORS: Heng Wang, Du Tran, Lorenzo Torresani, Matt Feiszli
HIGHLIGHT: This paper proposes an alternative approach based on a learnable correlation operator that can be used to establish frame-to-frame matches over convolutional feature maps in the different layers of the network.

37, TITLE: Projection & Probability-Driven Black-Box Attack
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Projection__Probability-Driven_Black-Box_Attack_CVPR_2020_paper.html
AUTHORS: Jie Li, Rongrong Ji, Hong Liu, Jianzhuang Liu, Bineng Zhong, Cheng Deng, Qi Tian
HIGHLIGHT: In this paper, we propose Projection & Probability-driven Black-box Attack (PPBA) to tackle this problem by reducing the solution space and providing better optimization.

38, TITLE: Auxiliary Training: Towards Accurate and Robust Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Auxiliary_Training_Towards_Accurate_and_Robust_Models_CVPR_2020_paper.html
AUTHORS: Linfeng Zhang, Muzhou Yu, Tong Chen, Zuoqiang Shi, Chenglong Bao, Kaisheng Ma
HIGHLIGHT: In this paper, we propose a novel training method via introducing the auxiliary classifiers for training on corrupted samples, while the clean samples are normally trained with the primary classifier.

39, TITLE: PaStaNet: Toward Human Activity Knowledge Engine
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_PaStaNet_Toward_Human_Activity_Knowledge_Engine_CVPR_2020_paper.html
AUTHORS: Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Shiyi Wang, Hao-Shu Fang, Ze Ma, Mingyang Chen, Cewu Lu
HIGHLIGHT: In light of this, we propose a new path: infer human part states first and then reason out the activities based on part-level semantics.

40, TITLE: A Hierarchical Graph Network for 3D Object Detection on Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_A_Hierarchical_Graph_Network_for_3D_Object_Detection_on_Point_CVPR_2020_paper.html
AUTHORS: Jintai Chen, Biwen Lei, Qingyu Song, Haochao Ying, Danny Z. Chen, Jian Wu
HIGHLIGHT: In this paper, we propose a new graph convolution (GConv) based hierarchical graph network (HGNet) for 3D object detection, which processes raw point clouds directly to predict 3D bounding boxes.

41, TITLE: Learning Generative Models of Shape Handles
http://openaccess.thecvf.com/content_CVPR_2020/html/Gadelha_Learning_Generative_Models_of_Shape_Handles_CVPR_2020_paper.html
AUTHORS: Matheus Gadelha, Giorgio Gori, Duygu Ceylan, Radomir Mech, Nathan Carr, Tamy Boubekeur, Rui Wang, Subhransu Maji
HIGHLIGHT: We present a generative model to synthesize 3D shapes as sets of handles -- lightweight proxies that approximate the original 3D shape -- for applications in interactive editing, shape parsing, and building compact 3D representations.

42, TITLE: One Man's Trash Is Another Man's Treasure: Resisting Adversarial Examples by Adversarial Examples
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiao_One_Mans_Trash_Is_Another_Mans_Treasure_Resisting_Adversarial_Examples_CVPR_2020_paper.html
AUTHORS: Chang Xiao, Changxi Zheng
HIGHLIGHT: We embrace the omnipresence of adversarial examples and the numerical procedure of crafting them, and turn this harmful attacking process into a useful defense mechanism.

43, TITLE: Toward a Universal Model for Shape From Texture
http://openaccess.thecvf.com/content_CVPR_2020/html/Verbin_Toward_a_Universal_Model_for_Shape_From_Texture_CVPR_2020_paper.html
AUTHORS: Dor Verbin, Todd Zickler
HIGHLIGHT: We consider the shape from texture problem, where the input is a single image of a curved, textured surface, and the texture and shape are both a priori unknown.

44, TITLE: HybridPose: 6D Object Pose Estimation Under Hybrid Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Song_HybridPose_6D_Object_Pose_Estimation_Under_Hybrid_Representations_CVPR_2020_paper.html
AUTHORS: Chen Song, Jiaru Song, Qixing Huang
HIGHLIGHT: We introduce HybridPose, a novel 6D object pose estimation approach.

45, TITLE: Boundary-Aware 3D Building Reconstruction From a Single Overhead Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Mahmud_Boundary-Aware_3D_Building_Reconstruction_From_a_Single_Overhead_Image_CVPR_2020_paper.html
AUTHORS: Jisan Mahmud, True Price, Akash Bapat, Jan-Michael Frahm
HIGHLIGHT: We propose a boundary-aware multi-task deep-learning-based framework for fast 3D building modeling from a single overhead image.

46, TITLE: Articulation-Aware Canonical Surface Mapping
http://openaccess.thecvf.com/content_CVPR_2020/html/Kulkarni_Articulation-Aware_Canonical_Surface_Mapping_CVPR_2020_paper.html
AUTHORS: Nilesh Kulkarni, Abhinav Gupta, David F. Fouhey, Shubham Tulsiani
HIGHLIGHT: We tackle the tasks of: 1) predicting a Canonical Surface Mapping (CSM) that indicates the mapping from 2D pixels to corresponding points on a canonical template shape , and 2) inferring the articulation and pose of the template corresponding to the input image.

47, TITLE: BiFuse: Monocular 360 Depth Estimation via Bi-Projection Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_BiFuse_Monocular_360_Depth_Estimation_via_Bi-Projection_Fusion_CVPR_2020_paper.html
AUTHORS: Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, Yi-Hsuan Tsai
HIGHLIGHT: Thus we propose a bi-projection fusion scheme along with learnable masks to balance the feature map from the two projections.

48, TITLE: Transformation GAN for Unsupervised Image Synthesis and Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Transformation_GAN_for_Unsupervised_Image_Synthesis_and_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Jiayu Wang, Wengang Zhou, Guo-Jun Qi, Zhongqian Fu, Qi Tian, Houqiang Li
HIGHLIGHT: To improve both image synthesis quality and representation learning performance under the unsupervised setting, in this paper, we propose a simple yet effective Transformation Generative Adversarial Networks (TrGAN).

49, TITLE: PPDM: Parallel Point Detection and Matching for Real-Time Human-Object Interaction Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Liao_PPDM_Parallel_Point_Detection_and_Matching_for_Real-Time_Human-Object_Interaction_CVPR_2020_paper.html
AUTHORS: Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, Jiashi Feng
HIGHLIGHT: We propose a single-stage Human-Object Interaction (HOI) detection method that has outperformed all existing methods on HICO-DET dataset at 37 fps on a single Titan XP GPU.

50, TITLE: Height and Uprightness Invariance for 3D Prediction From a Single View
http://openaccess.thecvf.com/content_CVPR_2020/html/Baradad_Height_and_Uprightness_Invariance_for_3D_Prediction_From_a_Single_CVPR_2020_paper.html
AUTHORS: Manel Baradad, Antonio Torralba
HIGHLIGHT: To account for this, we propose a system that directly regresses 3D world coordinates for each pixel.

51, TITLE: SCT: Set Constrained Temporal Transformer for Set Supervised Action Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Fayyaz_SCT_Set_Constrained_Temporal_Transformer_for_Set_Supervised_Action_Segmentation_CVPR_2020_paper.html
AUTHORS: Mohsen Fayyaz, Jurgen Gall
HIGHLIGHT: In order to address this task, we propose an approach that can be trained end-to-end on such data.

52, TITLE: 3DV: 3D Dynamic Voxel for Action Recognition in Depth Video
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_3DV_3D_Dynamic_Voxel_for_Action_Recognition_in_Depth_Video_CVPR_2020_paper.html
AUTHORS: Yancheng Wang, Yang Xiao, Fu Xiong, Wenxiang Jiang, Zhiguo Cao, Joey Tianyi Zhou, Junsong Yuan
HIGHLIGHT: With 3D space voxelization, the key idea of 3DV is to encode the 3D motion information within depth video into a regular voxel set (i.e., 3DV) compactly, via temporal rank pooling.

53, TITLE: Adaptive Interaction Modeling via Graph Operations Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Adaptive_Interaction_Modeling_via_Graph_Operations_Search_CVPR_2020_paper.html
AUTHORS: Haoxin Li, Wei-Shi Zheng, Yu Tao, Haifeng Hu, Jian-Huang Lai
HIGHLIGHT: In this paper, we automate the process of structures design to learn adaptive structures for interaction modeling.

54, TITLE:	Front2Back: Single View 3D Shape Reconstruction via Front to Back Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Yao_Front2Back_Single_View_3D_Shape_Reconstruction_via_Front_to_Back_CVPR_2020_paper.html
AUTHORS:	Yuan Yao, Nico Schertler, Enrique Rosales, Helge Rhodin, Leonid Sigal, Alla Sheffer
HIGHLIGHT:	In this work, we induce structure and geometric constraints by leveraging three core observations: (1) the surface of most everyday objects is often almost entirely exposed from pairs of typical opposite views; (2) everyday objects often exhibit global reflective symmetries which can be accurately predicted from single views; (3) opposite orthographic views of a 3D shape share consistent silhouettes.

55, TITLE:	SDC-Depth: Semantic Divide-and-Conquer Network for Monocular Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_SDC-Depth_Semantic_Divide-and-Conquer_Network_for_Monocular_Depth_Estimation_CVPR_2020_paper.html
AUTHORS:	Lijun Wang, Jianming Zhang, Oliver Wang, Zhe Lin, Huchuan Lu
HIGHLIGHT:	Due to its complexity, we propose a deep neural network model based on a semantic divide-and-conquer approach.

56, TITLE:	Single-View View Synthesis With Multiplane Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Tucker_Single-View_View_Synthesis_With_Multiplane_Images_CVPR_2020_paper.html
AUTHORS:	Richard Tucker, Noah Snavely
HIGHLIGHT:	Our method learns to predict a multiplane image directly from a single image input, and we introduce scale-invariant view synthesis for supervision, enabling us to train on online video.

57, TITLE:	Deep Parametric Shape Predictions Using Distance Fields
http://openaccess.thecvf.com/content_CVPR_2020/html/Smirnov_Deep_Parametric_Shape_Predictions_Using_Distance_Fields_CVPR_2020_paper.html
AUTHORS:	Dmitriy Smirnov, Matthew Fisher, Vladimir G. Kim, Richard Zhang, Justin Solomon
HIGHLIGHT:	Hence, we propose a new framework for predicting parametric shape primitives using deep learning.

58, TITLE:	Leveraging Photometric Consistency Over Time for Sparsely Supervised Hand-Object Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Hasson_Leveraging_Photometric_Consistency_Over_Time_for_Sparsely_Supervised_Hand-Object_Reconstruction_CVPR_2020_paper.html
AUTHORS:	Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, Cordelia Schmid
HIGHLIGHT:	To overcome this challenge we present a method to leverage photometric consistency across time when annotations are only available for a sparse subset of frames in a video.

59, TITLE:	Ensemble Generative Cleaning With Feedback Loops for Defending Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2020/html/Yuan_Ensemble_Generative_Cleaning_With_Feedback_Loops_for_Defending_Adversarial_Attacks_CVPR_2020_paper.html
AUTHORS:	Jianhe Yuan, Zhihai He
HIGHLIGHT:	In this paper, we develop a new method called ensemble generative cleaning with feedback loops (EGC-FL) for effective defense of deep neural networks.

60, TITLE:	Temporal Pyramid Network for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Temporal_Pyramid_Network_for_Action_Recognition_CVPR_2020_paper.html
AUTHORS:	Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, Bolei Zhou
HIGHLIGHT:	In this work we propose a generic Temporal Pyramid Network (TPN) at the feature-level, which can be flexibly integrated into 2D or 3D backbone networks in a plug-and-play manner.

61, TITLE:	FaceScape: A Large-Scale High Quality 3D Face Dataset and Detailed Riggable 3D Face Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_FaceScape_A_Large-Scale_High_Quality_3D_Face_Dataset_and_Detailed_CVPR_2020_paper.html
AUTHORS:	Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, Xun Cao
HIGHLIGHT:	In this paper, we present a large-scale detailed 3D face dataset, FaceScape, and propose a novel algorithm that is able to predict elaborate riggable 3D face models from a single image input.

62, TITLE:	Structure-Guided Ranking Loss for Single Image Depth Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Xian_Structure-Guided_Ranking_Loss_for_Single_Image_Depth_Prediction_CVPR_2020_paper.html

AUTHORS: Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, Zhiguo Cao
HIGHLIGHT: To more effectively learn from such pseudo-depth data, we propose to use a simple pair-wise ranking loss with a novel sampling strategy.

63, TITLE: In Perfect Shape: Certifiably Optimal 3D Shape Reconstruction From 2D Landmarks
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_In_Perfect_Shape_Certifiably_Optimal_3D_Shape_Reconstruction_From_2D_CVPR_2020_paper.html
AUTHORS: Heng Yang, Luca Carlone
HIGHLIGHT: We study the problem of 3D shape reconstruction from 2D landmarks extracted in a single image.

64, TITLE: When NAS Meets Robustness: In Search of Robust Architectures Against Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_When_NAS_Meets_Robustness_In_Search_of_Robust_Architectures_Against_CVPR_2020_paper.html
AUTHORS: Minghao Guo, Yuzhe Yang, Rui Xu, Ziwei Liu, Dahua Lin
HIGHLIGHT: In this work, we take an architectural perspective and investigate the patterns of network architectures that are resilient to adversarial attacks.

65, TITLE: Towards Transferable Targeted Attack
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Towards_Transferable_Targeted_Attack_CVPR_2020_paper.html
AUTHORS: Maosen Li, Cheng Deng, Tengjiao Li, Junchi Yan, Xinbo Gao, Heng Huang
HIGHLIGHT: To overcome the above problems, we propose a novel targeted attack approach to effectively generate more transferable adversarial examples.

66, TITLE: Self-Supervised Human Depth Estimation From Monocular Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Tan_Self-Supervised_Human_Depth_Estimation_From_Monocular_Videos_CVPR_2020_paper.html
AUTHORS: Feitong Tan, Hao Zhu, Zhaopeng Cui, Siyu Zhu, Marc Pollefeys, Ping Tan
HIGHLIGHT: This paper presents a self-supervised method that can be trained on YouTube videos without known depth, which makes training data collection simple and improves the generalization of the learned network.

67, TITLE: Recursive Social Behavior Graph for Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Recursive_Social_Behavior_Graph_for_Trajectory_Prediction_CVPR_2020_paper.html
AUTHORS: Jianhua Sun, Qinhong Jiang, Cewu Lu
HIGHLIGHT: In this paper, we present a novel insight of group-based social interaction model to explore relationships among pedestrians.

68, TITLE: Context-Aware and Scale-Insensitive Temporal Repetition Counting
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Context-Aware_and_Scale-Insensitive_Temporal_Repetition_Counting_CVPR_2020_paper.html
AUTHORS: Huaidong Zhang, Xuemiao Xu, Guoqiang Han, Shengfeng He
HIGHLIGHT: In this paper, we tailor a context-aware and scale-insensitive framework, to tackle the challenges in repetition counting caused by the unknown and diverse cycle-lengths.

69, TITLE: OASIS: A Large-Scale Dataset for Single Image 3D in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_OASIS_A_Large-Scale_Dataset_for_Single_Image_3D_in_the_CVPR_2020_paper.html
AUTHORS: Weifeng Chen, Shengyi Qian, David Fan, Noriyuki Kojima, Max Hamilton, Jia Deng
HIGHLIGHT: We address this issue by presenting Open Annotations of Single Image Surfaces (OASIS), a dataset for single-image 3D in the wild consisting of annotations of detailed 3D geometry for 140,000 images.

70, TITLE: VPLNet: Deep Single View Normal Estimation With Vanishing Points and Lines
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_VPLNet_Deep_Single_View_Normal_Estimation_With_Vanishing_Points_and_CVPR_2020_paper.html
AUTHORS: Rui Wang, David Geraghty, Kevin Matzen, Richard Szeliski, Jan-Michael Frahm
HIGHLIGHT: We present a novel single-view surface normal estimation method that combines traditional line and vanishing point analysis with a deep learning approach.

71, TITLE: Adversarial Robustness: From Self-Supervised Pre-Training to Fine-Tuning
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Adversarial_Robustness_From_Self-Supervised_Pre-Training_to_Fine-Tuning_CVPR_2020_paper.html

AUTHORS: Tianlong Chen, Sijia Liu, Shiyu Chang, Yu Cheng, Lisa Amini, Zhangyang Wang
HIGHLIGHT: We introduce adversarial training into self-supervision, to provide general-purpose robust pretrained models for the first time.

72, TITLE: Defending Against Universal Attacks Through Selective Feature Regeneration
http://openaccess.thecvf.com/content_CVPR_2020/html/Borkar_Defending_Against_Universal_Attacks_Through_Selective_Feature_Regeneration_CVPR_2020_paper.html
AUTHORS: Tejas Borkar, Felix Heide, Lina Karam
HIGHLIGHT: Departing from existing defense strategies that work mostly in the image domain, we present a novel defense which operates in the DNN feature domain and effectively defends against such universal perturbations.

73, TITLE: Universal Physical Camouflage Attacks on Object Detectors
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Universal_Physical_Camouflage_Attacks_on_Object_Detectors_CVPR_2020_paper.html
AUTHORS: Lifeng Huang, Chengying Gao, Yuyin Zhou, Cihang Xie, Alan L. Yuille, Changqing Zou, Ning Liu
HIGHLIGHT: In this paper, we study physical adversarial attacks on object detectors in the wild.

74, TITLE: Intra- and Inter-Action Understanding via Temporal Action Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Shao_Intra-_and_Inter-Action_Understanding_via_Temporal_Action_Parsing_CVPR_2020_paper.html
AUTHORS: Dian Shao, Yue Zhao, Bo Dai, Dahua Lin
HIGHLIGHT: Towards this goal, we construct TAPOS, a new dataset developed on sport videos with manual annotations of sub-actions, and conduct a study on temporal action parsing on top.

75, TITLE: Lightweight Photometric Stereo for Facial Details Recovery
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Lightweight_Photometric_Stereo_for_Facial_Details_Recovery_CVPR_2020_paper.html
AUTHORS: Xueying Wang, Yudong Guo, Bailin Deng, Juyong Zhang
HIGHLIGHT: In this paper, we present a lightweight strategy that only requires sparse inputs or even a single image to recover high-fidelity face shapes with images captured under near-field lights.

76, TITLE: Bundle Pooling for Polygonal Architecture Segmentation Problem
http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_Bundle_Pooling_for_Polygonal_Architecture_Segmentation_Problem_CVPR_2020_paper.html
AUTHORS: Huayi Zeng, Kevin Joseph, Adam Vest, Yasutaka Furukawa
HIGHLIGHT: This paper introduces a polygonal architecture segmentation problem, proposes bundle-pooling modules for line structure reasoning, and demonstrates a virtual remodeling application that produces production quality results.

77, TITLE: AvatarMe: Realistically Renderable 3D Facial Reconstruction "In-the-Wild"
http://openaccess.thecvf.com/content_CVPR_2020/html/Lattas_AvatarMe_Realistically_Renderable_3D_Facial_Reconstruction_In-the-Wild_CVPR_2020_paper.html
AUTHORS: Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we introduce AvatarMe, the first method that is able to reconstruct photorealistic 3D faces from a single "in-the-wild" image with an increasing level of detail.

78, TITLE: Defending Against Model Stealing Attacks With Adaptive Misinformation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kariyappa_Defending_Against_Model_Stealing_Attacks_With_Adaptive_Misinformation_CVPR_2020_paper.html
AUTHORS: Sanjay Kariyappa, Moinuddin K. Qureshi
HIGHLIGHT: We propose "Adaptive Misinformation" to defend against such model stealing attacks.

79, TITLE: Learning to Generate 3D Training Data Through Hybrid Gradient
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_to_Generate_3D_Training_Data_Through_Hybrid_Gradient_CVPR_2020_paper.html
AUTHORS: Dawei Yang, Jia Deng
HIGHLIGHT: In this work, we propose a new method that optimizes the generation of 3D training data based on what we call "hybrid gradient".

80, TITLE: Cascaded Refinement Network for Point Cloud Completion

http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Cascaded_Refinement_Network_for_Point_Cloud_Completion_CVPR_2020_paper.html
AUTHORS: Xiaogang Wang, Marcelo H. Ang Jr., Gim Hee Lee
HIGHLIGHT: To this end, we propose a cascaded refinement network together with a coarse-to-fine strategy to synthesize the detailed object shapes.

81, TITLE: Enhancing Intrinsic Adversarial Robustness via Feature Pyramid Decoder
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Enhancing_Intrinsic_Adversarial_Robustness_via_Feature_Pyramid_Decoder_CVPR_2020_paper.html
AUTHORS: Guanlin Li, Shuya Ding, Jun Luo, Chang Liu
HIGHLIGHT: In this paper, we propose an attack-agnostic defence framework to enhance the intrinsic robustness of neural networks, without jeopardizing the ability of generalizing clean samples.

82, TITLE: Learning to Discriminate Information for Online Action Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Eun_Learning_to_Discriminate_Information_for_Online_Action_Detection_CVPR_2020_paper.html
AUTHORS: Hyunjun Eun, Jinyoung Moon, Jongyoul Park, Chanho Jung, Changick Kim
HIGHLIGHT: For online action detection, in this paper, we propose a novel recurrent unit to explicitly discriminate the information relevant to an ongoing action from others.

83, TITLE: Adversarial Examples Improve Image Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_Adversarial_Examples_Improve_Image_Recognition_CVPR_2020_paper.html
AUTHORS: Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L. Yuille, Quoc V. Le
HIGHLIGHT: Here we present an opposite perspective: adversarial examples can be used to improve image recognition models if harnessed in the right manner.

84, TITLE: PQ-NET: A Generative Part Seq2Seq Network for 3D Shapes
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_PQ-NET_A_Generative_Part_Seq2Seq_Network_for_3D_Shapes_CVPR_2020_paper.html
AUTHORS: Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, Baoquan Chen
HIGHLIGHT: We introduce PQ-NET, a deep neural network which represents and generates 3D shapes via sequential part assembly.

85, TITLE: Actor-Transformers for Group Activity Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Gavrilyuk_Actor-Transformers_for_Group_Activity_Recognition_CVPR_2020_paper.html
AUTHORS: Kirill Gavrilyuk, Ryan Sanford, Mehrsan Javan, Cees G. M. Snoek
HIGHLIGHT: While existing solutions for this challenging problem explicitly model spatial and temporal relationships based on location of individual actors, we propose an actor-transformer model able to learn and selectively extract information relevant for group activity recognition.

86, TITLE: SG-NN: Sparse Generative Neural Networks for Self-Supervised Scene Completion of RGB-D Scans
http://openaccess.thecvf.com/content_CVPR_2020/html/Dai_SG-NN_Sparse_Generative_Neural_Networks_for_Self-Supervised_Scene_Completion_of_CVPR_2020_paper.html
AUTHORS: Angela Dai, Christian Diller, Matthias Niessner
HIGHLIGHT: We present a novel approach that converts partial and noisy RGB-D scans into high-quality 3D scene reconstructions by inferring unobserved scene geometry.

87, TITLE: Geometry-Aware Satellite-to-Ground Image Synthesis for Urban Areas
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Geometry-Aware_Satellite-to-Ground_Image_Synthesis_for_Urban_Areas_CVPR_2020_paper.html
AUTHORS: Xiaohu Lu, Zuoyue Li, Zhaopeng Cui, Martin R. Oswald, Marc Pollefeys, Rongjun Qin
HIGHLIGHT: We present a novel method for generating panoramic street-view images which are geometrically consistent with a given satellite image.

88, TITLE: Action Modifiers: Learning From Adverbs in Instructional Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Doughty_Action_Modifiers_Learning_From_Adverbs_in_Instructional_Videos_CVPR_2020_paper.html
AUTHORS: Hazel Doughty, Ivan Laptev, Walterio Mayol-Cuevas, Dima Damen
HIGHLIGHT: We present a method to learn a representation for adverbs from instructional videos using weak supervision from the accompanying narrations.

89, TITLE: ZSTAD: Zero-Shot Temporal Activity Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_ZSTAD_Zero-Shot_Temporal_Activity_Detection_CVPR_2020_paper.html
AUTHORS: Lingling Zhang, Xiaojun Chang, Jun Liu, Minnan Luo, Sen Wang, Zongyuan Ge, Alexander Hauptmann
HIGHLIGHT: To solve this challenging problem, we propose a novel task setting called zero-shot temporal activity detection (ZSTAD), where activities that have never been seen in training can still be detected.

90, TITLE: Geometric Structure Based and Regularized Depth Estimation From 360 Indoor Imagery
http://openaccess.thecvf.com/content_CVPR_2020/html/Jin_Geometric_Structure_Based_and_Regularized_Depth_Estimation_From_360_Indoor_CVPR_2020_paper.html
AUTHORS: Lei Jin, Yanyu Xu, Jia Zheng, Junfei Zhang, Rui Tang, Shugong Xu, Jingyi Yu, Shenghua Gao
HIGHLIGHT: Motivated by the correlation between the depth and the geometric structure of a 360 indoor image, we propose a novel learning-based depth estimation framework that leverages the geometric structure of a scene to conduct depth estimation.

91, TITLE: Deep Kinematics Analysis for Monocular 3D Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Deep_Kinematics_Analysis_for_Monocular_3D_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, Wenjun Zhang
HIGHLIGHT: In this paper, we propose to address above issue in a systematic view.

92, TITLE: TEA: Temporal Excitation and Aggregation for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_TEA_Temporal_Excitation_and_Aggregation_for_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, Limin Wang
HIGHLIGHT: In this paper, we propose a Temporal Excitation and Aggregation (TEA) block, including a motion excitation (ME) module and a multiple temporal aggregation (MTA) module, specifically designed to capture both short- and long-range temporal evolution.

93, TITLE: Oops! Predicting Unintentional Action in Video
http://openaccess.thecvf.com/content_CVPR_2020/html/Epstein_Oops_Predicting_Unintentional_Action_in_Video_CVPR_2020_paper.html
AUTHORS: Dave Epstein, Boyuan Chen, Carl Vondrick
HIGHLIGHT: We introduce a dataset of in-the-wild videos of unintentional action, as well as a suite of tasks for recognizing, localizing, and anticipating its onset.

94, TITLE: Scene Recomposition by Learning-Based ICP
http://openaccess.thecvf.com/content_CVPR_2020/html/Izadinia_Scene_Recomposition_by_Learning-Based_ICP_CVPR_2020_paper.html
AUTHORS: Hamid Izadinia, Steven M. Seitz
HIGHLIGHT: In addition to the fully automatic system, the key technical contribution is a novel approach for aligning CAD models to 3D scans, based on deep reinforcement learning.

95, TITLE: Enhancing Cross-Task Black-Box Transferability of Adversarial Examples With Dispersion Reduction
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Enhancing_Cross-Task_Black-Box_Transferability_of_Adversarial_Examples_With_Dispersion_Reduction_CVPR_2020_paper.html
AUTHORS: Yantao Lu, Yunhan Jia, Jianyu Wang, Bai Li, Weiheng Chai, Lawrence Carin, Senem Velipasalar
HIGHLIGHT: Our proposed attack minimizes the "dispersion" of the internal feature map, overcoming the limitations of existing attacks, that require task-specific loss functions and/or probing a target model.

96, TITLE: Deep Non-Line-of-Sight Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Chopite_Deep_Non-Line-of-Sight_Reconstruction_CVPR_2020_paper.html
AUTHORS: Javier Grau Chopite, Matthias B. Hullin, Michael Wand, Julian Iseringhausen
HIGHLIGHT: In this paper, we employ convolutional feed-forward networks for solving the reconstruction problem efficiently while maintaining good reconstruction quality.

97, TITLE: SSRNet: Scalable 3D Surface Reconstruction Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Mi_SSRNet_Scalable_3D_Surface_Reconstruction_Network_CVPR_2020_paper.html
AUTHORS: Zhenxing Mi, Yiming Luo, Wenbing Tao
HIGHLIGHT: In this paper, we propose the SSRNet, a novel scalable learning-based method for surface reconstruction.

98, TITLE: Progressive Relation Learning for Group Activity Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Progressive_Relation_Learning_for_Group_Activity_Recognition_CVPR_2020_paper.html
AUTHORS: Guyue Hu, Bo Cui, Yuan He, Shan Yu
HIGHLIGHT: In this paper, we propose a novel method based on deep reinforcement learning to progressively refine the low-level features and high-level relations of group activities.

99, TITLE: Cooling-Shrinking Attack: Blinding the Tracker With Imperceptible Noises
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Cooling-Shrinking_Attack_Blinding_the_Tracker_With_Imperceptible_Noises_CVPR_2020_paper.html
AUTHORS: Bin Yan, Dong Wang, Huchuan Lu, Xiaoyun Yang
HIGHLIGHT: In this paper, a cooling-shrinking attack method is proposed to deceive state-of-the-art SiameseRPN-based trackers.

100, TITLE: Adversarial Camouflage: Hiding Physical-World Attacks With Natural Styles
http://openaccess.thecvf.com/content_CVPR_2020/html/Duan_Adversarial_Camouflage_Hiding_Physical-World_Attacks_With_Natural_Styles_CVPR_2020_paper.html
AUTHORS: Ranjie Duan, Xingjun Ma, Yisen Wang, James Bailey, A. K. Qin, Yun Yang
HIGHLIGHT: In this paper, we propose a novel approach, called Adversarial Camouflage (AdvCam), to craft and camouflage physical-world adversarial examples into natural styles that appear legitimate to human observers.

101, TITLE: Weakly-Supervised Action Localization by Generative Attention Modeling
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Weakly-Supervised_Action_Localization_by_Generative_Attention_Modeling_CVPR_2020_paper.html
AUTHORS: Baifeng Shi, Qi Dai, Yadong Mu, Jingdong Wang
HIGHLIGHT: To solve the problem, in this paper we propose to model the class-agnostic frame-wise probability conditioned on the frame attention using conditional Variational Auto-Encoder (VAE).

102, TITLE: Towards Achieving Adversarial Robustness by Enforcing Feature Consistency Across Bit Planes
http://openaccess.thecvf.com/content_CVPR_2020/html/Addepalli_Towards_Achieving_Adversarial_Robustness_by_Enforcing_Feature_Consistency_Across_Bit_CVPR_2020_paper.html
AUTHORS: Sravanti Addepalli, Vivek B.S., Arya Baburaj, Gaurang Sriramanan, R. Venkatesh Babu
HIGHLIGHT: In this work, we attempt to address this problem by training networks to form coarse impressions based on the information in higher bit planes, and use the lower bit planes only to refine their prediction.

103, TITLE: Polishing Decision-Based Adversarial Noise With a Customized Sampling
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Polishing_Decision-Based_Adversarial_Noise_With_a_Customized_Sampling_CVPR_2020_paper.html
AUTHORS: Yucheng Shi, Yahong Han, Qi Tian
HIGHLIGHT: In this paper, we demonstrate the advantage of using current noise and historical queries to customize the variance and mean of sampling in boundary attack to polish adversarial noise.

104, TITLE: Towards Large Yet Imperceptible Adversarial Image Perturbations With Perceptual Color Distance
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Towards_Large_Yet_Imperceptible_Adversarial_Image_Perturbations_With_Perceptual_Color_CVPR_2020_paper.html
AUTHORS: Zhengyu Zhao, Zhuoran Liu, Martha Larson
HIGHLIGHT: In this work, we drop this assumption by pursuing an approach that exploits human color perception, and more specifically, minimizing perturbation size with respect to perceptual color distance.

105, TITLE: Something-Else: Compositional Action Recognition With Spatial-Temporal Interaction Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Materzynska_Something-Else_Compositional_Action_Recognition_With_Spatial-Temporal_Interaction_Networks_CVPR_2020_paper.html
AUTHORS: Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, Trevor Darrell
HIGHLIGHT: In this paper, we study the compositionality of action by looking into the dynamics of subject-object interactions.

106, TITLE: Learning Unsupervised Hierarchical Part Decomposition of 3D Objects From a Single RGB Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Paschalidou_Learning_Unsupervised_Hierarchical_Part_Decomposition_of_3D_Objects_From_a_CVPR_2020_paper.html
AUTHORS: Despoina Paschalidou, Luc Van Gool, Andreas Geiger

HIGHLIGHT: We address this challenging problem by proposing a novel formulation that allows to jointly recover the geometry of a 3D object as a set of primitives as well as their latent hierarchical structure without part-level supervision.

107, TITLE: Focus on Defocus: Bridging the Synthetic to Real Domain Gap for Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Maximov_Focus_on_Defocus_Bridging_the_Synthetic_to_Real_Domain_Gap_CVPR_2020_paper.html
AUTHORS: Maxim Maximov, Kevin Galim, Laura Leal-Taixe
HIGHLIGHT: In this paper, we tackle this issue by using domain invariant defocus blur as direct supervision.

108, TITLE: Active Vision for Early Recognition of Human Actions
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Active_Vision_for_Early_Recognition_of_Human_Actions_CVPR_2020_paper.html
AUTHORS: Boyu Wang, Lihan Huang, Minh Hoai
HIGHLIGHT: We propose a method for early recognition of human actions, one that can take advantages of multiple cameras while satisfying the constraints due to limited communication bandwidth and processing power.

109, TITLE: SmallBigNet: Integrating Core and Contextual Views for Video Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_SmallBigNet_Integrating_Core_and_Contextual_Views_for_Video_Classification_CVPR_2020_paper.html
AUTHORS: Xianhang Li, Yali Wang, Zhipeng Zhou, Yu Qiao
HIGHLIGHT: To alleviate this problem, we propose a concise and novel SmallBig network, with the cooperation of small and big views.

110, TITLE: Gate-Shift Networks for Video Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Sudhakaran_Gate-Shift_Networks_for_Video_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Swathikiran Sudhakaran, Sergio Escalera, Oswald Lanz
HIGHLIGHT: In this paper we introduce spatial gating in spatial-temporal decomposition of 3D kernels.

111, TITLE: Semantics-Guided Neural Networks for Efficient Skeleton-Based Human Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Semantics-Guided_Neural_Networks_for_Efficient_Skeleton-Based_Human_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Pengfei Zhang, Cuiling Lan, Wenjun Zeng, Junliang Xing, Jianru Xue, Nanning Zheng
HIGHLIGHT: In this paper, we propose a simple yet effective semantics-guided neural network (SGN) for skeleton-based action recognition.

112, TITLE: Exploiting Joint Robustness to Adversarial Perturbations
http://openaccess.thecvf.com/content_CVPR_2020/html/Dabouei_Exploiting_Joint_Robustness_to_Adversarial_Perturbations_CVPR_2020_paper.html
AUTHORS: Ali Dabouei, Sobhan Soleymani, Fariborz Taherkhani, Jeremy Dawson, Nasser M. Nasrabadi
HIGHLIGHT: In this paper, we exploit first-order interactions within ensembles to formalize a reliable and practical defense.

113, TITLE: From Image Collections to Point Clouds With Self-Supervised Shape and Pose Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Navaneet_From_Image_Collections_to_Point_Clouds_With_Self-Supervised_Shape_and_CVPR_2020_paper.html
AUTHORS: K L Navaneet, Ansu Mathew, Shashank Kashyap, Wei-Chih Hung, Varun Jampani, R. Venkatesh Babu
HIGHLIGHT: In this work, we propose a deep learning technique for 3D object reconstruction from a single image.

114, TITLE: Searching for Actions on the Hyperbole
http://openaccess.thecvf.com/content_CVPR_2020/html/Long_Searching_for_Actions_on_the_Hyperbole_CVPR_2020_paper.html
AUTHORS: Teng Long, Pascal Mettes, Heng Tao Shen, Cees G. M. Snoek
HIGHLIGHT: In this paper, we introduce hierarchical action search.

115, TITLE: ColorFool: Semantic Adversarial Colorization
http://openaccess.thecvf.com/content_CVPR_2020/html/Shamsabadi_ColorFool_Semantic_Adversarial_Colorization_CVPR_2020_paper.html
AUTHORS: Ali Shahin Shamsabadi, Ricardo Sanchez-Matilla, Andrea Cavallaro
HIGHLIGHT: In this paper, we propose a content-based black-box adversarial attack that generates unrestricted perturbations by exploiting image semantics to selectively modify colors within chosen ranges that are perceived as natural by humans.

116, TITLE: Boosting the Transferability of Adversarial Samples via Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Boosting_the_Transferability_of_Adversarial_Samples_via_Attention_CVPR_2020_paper.html
AUTHORS: Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R. Lyu, Yu-Wing Tai
HIGHLIGHT: In this work, we propose a novel mechanism to alleviate the overfitting issue.

117, TITLE: ActionBytes: Learning From Trimmed Videos to Localize Actions
http://openaccess.thecvf.com/content_CVPR_2020/html/Jain_ActionBytes_Learning_From_Trimmed_Videos_to_Localize_Actions_CVPR_2020_paper.html
AUTHORS: Mihir Jain, Amir Ghodrati, Cees G. M. Snoek
HIGHLIGHT: We propose a method to train an action localization network that segments a video into interpretable fragments, we call ActionBytes.

118, TITLE: Efficient Adversarial Training With Transferable Adversarial Examples
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Efficient_Adversarial_Training_With_Transferable_Adversarial_Examples_CVPR_2020_paper.html
AUTHORS: Haizhong Zheng, Ziqi Zhang, Juncheng Gu, Honglak Lee, Atul Prakash
HIGHLIGHT: Leveraging this property, we propose a novel method, Adversarial Training with Transferable Adversarial Examples (ATTA), that can enhance the robustness of trained models and greatly improve the training efficiency by accumulating adversarial perturbations through epochs.

119, TITLE: Alleviation of Gradient Exploding in GANs: Fake Can Be Real
http://openaccess.thecvf.com/content_CVPR_2020/html/Tao_Alleviation_of_Gradient_Exploding_in_GANs_Fake_Can_Be_Real_CVPR_2020_paper.html
AUTHORS: Song Tao, Jia Wang
HIGHLIGHT: In order to alleviate the notorious mode collapse phenomenon in generative adversarial networks (GANs), we propose a novel training method of GANs in which certain fake samples are considered as real ones during the training process.

120, TITLE: On Isometry Robustness of Deep 3D Point Cloud Models Under Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_On_Isometry_Robustness_of_Deep_3D_Point_Cloud_Models_Under_CVPR_2020_paper.html
AUTHORS: Yue Zhao, Yuwei Wu, Caihua Chen, Andrew Lim
HIGHLIGHT: Incorporating with the Restricted Isometry Property, we propose a novel framework of white-box attack on top of spectral norm based perturbation.

121, TITLE: Achieving Robustness in the Wild via Adversarial Mixing With Disentangled Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Gowal_Achieving_Robustness_in_the_Wild_via_Adversarial_Mixing_With_Disentangled_CVPR_2020_paper.html
AUTHORS: Sven Gowal, Chongli Qin, Po-Sen Huang, Taylan Cemgil, Krishnamurthy Dvijotham, Timothy Mann, Pushmeet Kohli
HIGHLIGHT: In this paper, we propose a novel approach to express and formalize robustness to these kinds of real-world transformations of the input.

122, TITLE: QEBA: Query-Efficient Boundary-Based Blackbox Attack
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_QEBA_Query-Efficient_Boundary-Based_Blackbox_Attack_CVPR_2020_paper.html
AUTHORS: Huichen Li, Xiaojun Xu, Xiaolu Zhang, Shuang Yang, Bo Li
HIGHLIGHT: In this paper, we propose a Query-Efficient Boundary-based blackbox Attack (QEBA) based only on model's final prediction labels.

123, TITLE: Learning to Simulate Dynamic Environments With GameGAN
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Learning_to_Simulate_Dynamic_Environments_With_GameGAN_CVPR_2020_paper.html
AUTHORS: Seung Wook Kim, Yuhao Zhou, Jonah Philion, Antonio Torralba, Sanja Fidler
HIGHLIGHT: In this paper, we aim to learn a simulator by simply watching an agent interact with an environment.

124, TITLE: Learn2Perturb: An End-to-End Feature Perturbation Learning to Improve Adversarial Robustness
http://openaccess.thecvf.com/content_CVPR_2020/html/Jeddi_Learn2Perturb_An_End-to-End_Feature_Perturbation_Learning_to_Improve_Adversarial_Robustness_CVPR_2020_paper.html
AUTHORS: Ahmadreza Jeddi, Mohammad Javad Shafiee, Michelle Karg, Christian Scharfenberger, Alexander Wong
HIGHLIGHT: In this study, we introduce Learn2Perturb, an end-to-end feature perturbation learning approach for improving the adversarial robustness of deep neural networks.

125, TITLE:	SDFDiff: Differentiable Rendering of Signed Distance Fields for 3D Shape Optimization
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_SDFDiff_Differentiable_Rendering_of_Signed_Distance_Fields_for_3D_Shape_CVPR_2020_paper.html
AUTHORS:	Yue Jiang, Dantong Ji, Zhizhong Han, Matthias Zwicker
HIGHLIGHT:	We propose SDFDiff, a novel approach for image-based shape optimization using differentiable rendering of 3D shapes represented by signed distance functions (SDFs).

126, TITLE:	Through the Looking Glass: Neural 3D Reconstruction of Transparent Shapes
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Through_the_Looking_Glass_Neural_3D_Reconstruction_of_Transparent_Shapes_CVPR_2020_paper.html
AUTHORS:	Zhengqin Li, Yu-Ying Yeh, Manmohan Chandraker
HIGHLIGHT:	Our novel contributions include a normal representation that enables the network to model complex light transport through local computation, a rendering layer that models refractions and reflections, a cost volume specifically designed for normal refinement of transparent shapes and a feature mapping based on predicted normals for 3D point cloud reconstruction.

127, TITLE:	TextureFusion: High-Quality Texture Acquisition for Real-Time RGB-D Scanning
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_TextureFusion_High-Quality_Texture_Acquisition_for_Real-Time_RGB-D_Scanning_CVPR_2020_paper.html
AUTHORS:	Joo Ho Lee, Hyunho Ha, Yue Dong, Xin Tong, Min H. Kim
HIGHLIGHT:	In this work, we propose a progressive texture-fusion method specially designed for real-time RGB-D scanning.

128, TITLE:	D3VO: Deep Depth, Deep Pose and Deep Uncertainty for Monocular Visual Odometry
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_D3VO_Deep_Depth_Deep_Pose_and_Deep_Uncertainty_for_Monocular_CVPR_2020_paper.html
AUTHORS:	Nan Yang, Lukas von Stumberg, Rui Wang, Daniel Cremers
HIGHLIGHT:	We propose D3VO as a novel framework for monocular visual odometry that exploits deep networks on three levels -- deep depth, pose and uncertainty estimation.

129, TITLE:	Deep Implicit Volume Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Deep_Implicit_Volume_Compression_CVPR_2020_paper.html
AUTHORS:	Danhang Tang, Saurabh Singh, Philip A. Chou, Christian Hane, Mingsong Dou, Sean Fanello, Jonathan Taylor, Philip Davidson, Onur G. Guleryuz, Yinda Zhang, Shahram Izadi, Andrea Tagliasacchi, Sofien Bouaziz, Cem Keskin
HIGHLIGHT:	We describe a novel approach for compressing truncated signed distance fields (TSDF) stored in 3D voxel grids, and their corresponding textures.

130, TITLE:	MAGSAC++, a Fast, Reliable and Accurate Robust Estimator
http://openaccess.thecvf.com/content_CVPR_2020/html/Barath_MAGSAC_a_Fast_Reliable_and_Accurate_Robust_Estimator_CVPR_2020_paper.html
AUTHORS:	Daniel Barath, Jana Noskova, Maksym Ivashechkin, Jiri Matas
HIGHLIGHT:	We propose MAGSAC++ and Progressive NAPSAC sampler, P-NAPSAC in short.

131, TITLE:	OctSqueeze: Octree-Structured Entropy Model for LiDAR Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_OctSqueeze_Octree-Structured_Entropy_Model_for_LiDAR_Compression_CVPR_2020_paper.html
AUTHORS:	Lila Huang, Shenlong Wang, Kelvin Wong, Jerry Liu, Raquel Urtasun
HIGHLIGHT:	We present a novel deep compression algorithm to reduce the memory footprint of LiDAR point clouds.

132, TITLE:	4D Association Graph for Realtime Multi-Person Motion Capture Using Multiple Video Cameras
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_4D_Association_Graph_for_Realtime_Multi-Person_Motion_Capture_Using_Multiple_CVPR_2020_paper.html
AUTHORS:	Yuxiang Zhang, Liang An, Tao Yu, Xiu Li, Kun Li, Yebin Liu
HIGHLIGHT:	This paper contributes a novel realtime multi-person motion capture algorithm using multiview video inputs.

133, TITLE:	Upgrading Optical Flow to 3D Scene Flow Through Optical Expansion
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Upgrading_Optical_Flow_to_3D_Scene_Flow_Through_Optical_Expansion_CVPR_2020_paper.html
AUTHORS:	Gengshan Yang, Deva Ramanan
HIGHLIGHT:	We describe an approach for upgrading 2D optical flow to 3D scene flow.

134, TITLE: Robust 3D Self-Portraits in Seconds
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Robust_3D_Self-Portraits_in_Seconds_CVPR_2020_paper.html
AUTHORS: Zhe Li, Tao Yu, Chuanyu Pan, Zerong Zheng, Yebin Liu
HIGHLIGHT: In this paper, we propose an efficient method for robust 3D self-portraits using a single RGBD camera.

135, TITLE: FastDVDnet: Towards Real-Time Deep Video Denoising Without Flow Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Tassano_FastDVDnet_Towards_Real-Time_Deep_Video_Denoising_Without_Flow_Estimation_CVPR_2020_paper.html
AUTHORS: Matias Tassano, Julie Delon, Thomas Veit
HIGHLIGHT: In this paper, we propose a state-of-the-art video denoising algorithm based on a convolutional neural network architecture.

136, TITLE: Learning to Have an Ear for Face Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Meishvili_Learning_to_Have_an_Ear_for_Face_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Givi Meishvili, Simon Jenni, Paolo Favaro
HIGHLIGHT: We propose a novel method to use both audio and a low-resolution image to perform extreme face super-resolution (a 16x increase of the input size).

137, TITLE: Deep Optics for Single-Shot High-Dynamic-Range Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Metzler_Deep_Optics_for_Single-Shot_High-Dynamic-Range_Imaging_CVPR_2020_paper.html
AUTHORS: Christopher A. Metzler, Hayato Ikoma, Yifan Peng, Gordon Wetzstein
HIGHLIGHT: Inspired by recent deep optical imaging approaches, we interpret this problem as jointly training an optical encoder and electronic decoder where the encoder is parameterized by the point spread function (PSF) of the lens, the bottleneck is the sensor with a limited dynamic range, and the decoder is a convolutional neural network (CNN).

138, TITLE: Learning Rank-1 Diffractive Optics for Single-Shot High Dynamic Range Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Learning_Rank-1_Diffractive_Optics_for_Single-Shot_High_Dynamic_Range_Imaging_CVPR_2020_paper.html
AUTHORS: Qilin Sun, Ethan Tseng, Qiang Fu, Wolfgang Heidrich, Felix Heide
HIGHLIGHT: In this work, we propose a method for snapshot HDR imaging by learning an optical HDR encoding in a single image which maps saturated highlights into neighboring unsaturated areas using a diffractive optical element (DOE).

139, TITLE: Deep White-Balance Editing
http://openaccess.thecvf.com/content_CVPR_2020/html/Afifi_Deep_White-Balance_Editing_CVPR_2020_paper.html
AUTHORS: Mahmoud Afifi, Michael S. Brown
HIGHLIGHT: We introduce a deep learning approach to realistically edit an sRGB image's white balance.

140, TITLE: Non-Line-of-Sight Surface Reconstruction Using the Directional Light-Cone Transform
http://openaccess.thecvf.com/content_CVPR_2020/html/Young_Non-Line-of-Sight_Surface_Reconstruction_Using_the_Directional_Light-Cone_Transform_CVPR_2020_paper.html
AUTHORS: Sean I. Young, David B. Lindell, Bernd Girod, David Taubman, Gordon Wetzstein
HIGHLIGHT: We propose a joint albedo-normal approach to non-line-of-sight (NLOS) surface reconstruction using the directional light-cone transform (D-LCT).

141, TITLE: Seeing the World in a Bag of Chips
http://openaccess.thecvf.com/content_CVPR_2020/html/Park_Seeing_the_World_in_a_Bag_of_Chips_CVPR_2020_paper.html
AUTHORS: Jeong Joon Park, Aleksander Holynski, Steven M. Seitz
HIGHLIGHT: Our contributions include 1) modeling highly specular objects, 2) modeling inter-reflections and Fresnel effects, and 3) enabling surface light field reconstruction with the same input needed to reconstruct shape alone.

142, TITLE: Correction Filter for Single Image Super-Resolution: Robustifying Off-the-Shelf Deep Super-Resolvers
http://openaccess.thecvf.com/content_CVPR_2020/html/Abu_Hussein_Correction_Filter_for_Single_Image_Super-Resolution_Robustifying_Off-the-Shelf_Deep_Super-Resolvers_CVPR_2020_paper.html
AUTHORS: Shady Abu Hussein, Tom Tirer, Raja Giryes
HIGHLIGHT: Inspired by the literature on generalized sampling, in this work we propose a method for improving the performance of DNNs that have been trained with a fixed kernel on observations acquired by other kernels.

143, TITLE: Retina-Like Visual Image Reconstruction via Spiking Neural Model

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Retina-
Like_Visual_Image_Reconstruction_via_Spiking_Neural_Model_CVPR_2020_paper.html
AUTHORS: Lin Zhu, Siwei Dong, Jianing Li, Tiejun Huang, Yonghong Tian
HIGHLIGHT: In this paper, we design a retina-like visual image reconstruction framework, which is flexible in reconstructing full texture of natural scenes from the totally new spike data.

144, TITLE: Plug-and-Play Algorithms for Large-Scale Snapshot Compressive Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Yuan_Plug-and-Play_Algorithms_for_Large-
Scale_Snapshot_Compressive_Imaging_CVPR_2020_paper.html
AUTHORS: Xin Yuan, Yang Liu, Jinli Suo, Qionghai Dai
HIGHLIGHT: In this paper, we develop fast and flexible algorithms for SCI based on the plug-and-play (PnP) framework.

145, TITLE: Neural Network Pruning With Residual-Connections and Limited-Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Neural_Network_Pruning_With_Residual-Connections_and_Limited-
Data_CVPR_2020_paper.html
AUTHORS: Jian-Hao Luo, Jianxin Wu
HIGHLIGHT: In order to avoid the influence of label noise, we propose a label refinement approach to solve this problem.

146, TITLE: AdderNet: Do We Really Need Multiplications in Deep Learning?
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_AdderNet_Do_We_Really_Need_Multiplications_in_Deep_Learning_
CVPR_2020_paper.html
AUTHORS: Hanting Chen, Yunhe Wang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, Chang Xu
HIGHLIGHT: In this paper, we present adder networks (AdderNets) to trade these massive multiplications in deep neural networks, especially convolutional neural networks (CNNs), for much cheaper additions to reduce computation costs.

147, TITLE: NeuralScale: Efficient Scaling of Neurons for Resource-Constrained Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_NeuralScale_Efficient_Scaling_of_Neurons_for_Resource-
Constrained_Deep_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Eugene Lee, Chen-Yi Lee
HIGHLIGHT: In this work, we attempt to search for the neuron (filter) configuration of a fixed network architecture that maximizes accuracy.

148, TITLE: Training Quantized Neural Networks With a Full-Precision Auxiliary Module
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhuang_Training_Quantized_Neural_Networks_With_a_Full-
Precision_Auxiliary_Module_CVPR_2020_paper.html
AUTHORS: Bohan Zhuang, Lingqiao Liu, Mingkui Tan, Chunhua Shen, Ian Reid
HIGHLIGHT: In this paper, we seek to tackle a challenge in training low-precision networks: the notorious difficulty in propagating gradient through a low-precision network due to the non-differentiable quantization function.

149, TITLE: Neural Networks Are More Productive Teachers Than Human Raters: Active Mixup for Data-Efficient Knowledge Distillation From a Blackbox Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Neural_Networks_Are_More_Productive_Teachers_Than_Human_R
aters_Active_CVPR_2020_paper.html
AUTHORS: Dongdong Wang, Yandong Li, Liqiang Wang, Boqing Gong
HIGHLIGHT: To tackle these challenges, we propose an approach that blends mixup and active learning.

150, TITLE: Multi-Dimensional Pruning: A Unified Framework for Model Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Multi-
Dimensional_Pruning_A_Unified_Framework_for_Model_Compression_CVPR_2020_paper.html
AUTHORS: Jinyang Guo, Wanli Ouyang, Dong Xu
HIGHLIGHT: In this work, we propose a unified model compression framework called Multi-Dimensional Pruning (MDP) to simultaneously compress the convolutional neural networks (CNNs) on multiple dimensions.

151, TITLE: Towards Efficient Model Compression via Learned Global Ranking
http://openaccess.thecvf.com/content_CVPR_2020/html/Chin_Towards_Efficient_Model_Compression_via_Learned_Global_Rankin
g_CVPR_2020_paper.html
AUTHORS: Ting-Wu Chin, Ruizhou Ding, Cha Zhang, Diana Marculescu
HIGHLIGHT: To this end, we propose to learn a global ranking of the filters across different layers of the ConvNet, which is used to obtain a set of ConvNet architectures that have different accuracy/latency trade-offs by pruning the bottom-ranked filters.

152, TITLE: HRank: Filter Pruning Using High-Rank Feature Map

17

http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_HRank_Filter_Pruning_Using_High-Rank_Feature_Map_CVPR_2020_paper.html
AUTHORS: Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, Ling Shao
HIGHLIGHT: In this paper, we propose a novel filter pruning method by exploring the High Rank of feature maps (HRank).

153, TITLE: DMCP: Differentiable Markov Channel Pruning for Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_DMCP_Differentiable_Markov_Channel_Pruning_for_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Shaopeng Guo, Yujie Wang, Quanquan Li, Junjie Yan
HIGHLIGHT: In this paper, we propose a novel differentiable method for channel pruning, named Differentiable Markov Channel Pruning (DMCP), to efficiently search the optimal sub-structure.

154, TITLE: ReSprop: Reuse Sparsified Backpropagation
http://openaccess.thecvf.com/content_CVPR_2020/html/Goli_ReSprop_Reuse_Sparsified_Backpropagation_CVPR_2020_paper.html
AUTHORS: Negar Goli, Tor M. Aamodt
HIGHLIGHT: In this work, we focus on accelerating training by observing that about 90% of gradients are reusable during training.

155, TITLE: Adversarial Texture Optimization From RGB-D Scans
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Adversarial_Texture_Optimization_From_RGB-D_Scans_CVPR_2020_paper.html
AUTHORS: Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu "Max" Jiang, Leonidas J. Guibas, Matthias Niessner, Thomas Funkhouser
HIGHLIGHT: In this work, we present a novel approach for color texture generation using a conditional adversarial loss obtained from weakly-supervised views.

156, TITLE: Synchronizing Probability Measures on Rotations via Optimal Transport
http://openaccess.thecvf.com/content_CVPR_2020/html/Birdal_Synchronizing_Probability_Measures_on_Rotations_via_Optimal_Transport_CVPR_2020_paper.html
AUTHORS: Tolga Birdal, Michael Arbel, Umut Simsekli, Leonidas J. Guibas
HIGHLIGHT: We propose a nonparametric Riemannian particle optimization approach to solve the problem.

157, TITLE: GhostNet: More Features From Cheap Operations
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_GhostNet_More_Features_From_Cheap_Operations_CVPR_2020_paper.html
AUTHORS: Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, Chang Xu
HIGHLIGHT: This paper proposes a novel Ghost module to generate more feature maps from cheap operations.

158, TITLE: Attention-Aware Multi-View Stereo
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Attention-Aware_Multi-View_Stereo_CVPR_2020_paper.html
AUTHORS: Keyang Luo, Tao Guan, Lili Ju, Yuesong Wang, Zhuo Chen, Yawei Luo
HIGHLIGHT: In this paper, we propose an attention-aware deep neural network "AttMVS" for learning multi-view stereo.

159, TITLE: Bi3D: Stereo Depth Estimation via Binary Classifications
http://openaccess.thecvf.com/content_CVPR_2020/html/Badki_Bi3D_Stereo_Depth_Estimation_via_Binary_Classifications_CVPR_2020_paper.html
AUTHORS: Abhishek Badki, Alejandro Troccoli, Kihwan Kim, Jan Kautz, Pradeep Sen, Orazio Gallo
HIGHLIGHT: We present Bi3D, a method that estimates depth via a series of binary classifications.

160, TITLE: Joint Filtering of Intensity Images and Neuromorphic Events for High-Resolution Noise-Robust Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Joint_Filtering_of_Intensity_Images_and_Neuromorphic_Events_for_High-Resolution_CVPR_2020_paper.html
AUTHORS: Zihao W. Wang, Peiqi Duan, Oliver Cossairt, Aggelos Katsaggelos, Tiejun Huang, Boxin Shi
HIGHLIGHT: We present a novel computational imaging system with high resolution and low noise.

161, TITLE: SGAS: Sequential Greedy Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_SGAS_Sequential_Greedy_Architecture_Search_CVPR_2020_paper.html
AUTHORS: Guohao Li, Guocheng Qian, Itzel C. Delgadillo, Matthias Muller, Ali Thabet, Bernard Ghanem

HIGHLIGHT: Aiming to alleviate this common issue, we introduce sequential greedy architecture search (SGAS), an efficient method for neural architecture search.

162, TITLE: HVNet: Hybrid Voxel Network for LiDAR Based 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_HVNet_Hybrid_Voxel_Network_for_LiDAR_Based_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Maosheng Ye, Shuangjie Xu, Tongyi Cao
HIGHLIGHT: We present Hybrid Voxel Network (HVNet), a novel one-stage unified network for point cloud based 3D object detection for autonomous driving.

163, TITLE: Frequency Domain Compact 3D Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Frequency_Domain_Compact_3D_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Hanting Chen, Yunhe Wang, Han Shu, Yehui Tang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, Chang Xu
HIGHLIGHT: In this paper, we develop a novel approach for eliminating redundancy in the time dimensionality of 3D convolution filters by converting them into the frequency domain through a series of learned optimal transforms with extremely fewer parameters.

164, TITLE: Single-Image HDR Reconstruction by Learning to Reverse the Camera Pipeline
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Single-Image_HDR_Reconstruction_by_Learning_to_Reverse_the_Camera_Pipeline_CVPR_2020_paper.html
AUTHORS: Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, Jia-Bin Huang
HIGHLIGHT: In contrast to existing learning-based methods, our core idea is to incorporate the domain knowledge of the LDR image formation pipeline into our model.

165, TITLE: DNU: Deep Non-Local Unrolling for Computational Spectral Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_DNU_Deep_Non-Local_Unrolling_for_Computational_Spectral_Imaging_CVPR_2020_paper.html
AUTHORS: Lizhi Wang, Chen Sun, Maoqing Zhang, Ying Fu, Hua Huang
HIGHLIGHT: In this paper, we propose an interpretable neural network for computational spectral imaging.

166, TITLE: Single Image Optical Flow Estimation With an Event Camera
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Single_Image_Optical_Flow_Estimation_With_an_Event_Camera_CVPR_2020_paper.html
AUTHORS: Liyuan Pan, Miaomiao Liu, Richard Hartley
HIGHLIGHT: In this paper, we propose a single image (potentially blurred) and events based optical flow estimation approach.

167, TITLE: Multi-View Neural Human Rendering
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Multi-View_Neural_Human_Rendering_CVPR_2020_paper.html
AUTHORS: Minye Wu, Yuehao Wang, Qiang Hu, Jingyi Yu
HIGHLIGHT: We present an end-to-end Neural Human Renderer (NHR) for dynamic human captures under the multi-view setting.

168, TITLE: Depth Sensing Beyond LiDAR Range
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Depth_Sensing_Beyond_LiDAR_Range_CVPR_2020_paper.html
AUTHORS: Kai Zhang, Jiaxin Xie, Noah Snavely, Qifeng Chen
HIGHLIGHT: To that end, we propose a novel three-camera system that utilizes small field of view cameras.

169, TITLE: Event Probability Mask (EPM) and Event Denoising Convolutional Neural Network (EDnCNN) for Neuromorphic Cameras
http://openaccess.thecvf.com/content_CVPR_2020/html/Baldwin_Event_Probability_Mask_EPM_and_Event_Denoising_Convolutional_Neural_Network_CVPR_2020_paper.html
AUTHORS: R. Wes Baldwin, Mohammed Almatrafi, Vijayan Asari, Keigo Hirakawa
HIGHLIGHT: This paper presents a novel method for labeling real-world neuromorphic camera sensor data by calculating the likelihood of generating an event at each pixel within a short time window, which we refer to as "event probability mask" or EPM.

170, TITLE: Point-GNN: Graph Neural Network for 3D Object Detection in a Point Cloud

http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Point-
GNN_Graph_Neural_Network_for_3D_Object_Detection_in_a_CVPR_2020_paper.html
AUTHORS: Weijing Shi, Raj Rajkumar
HIGHLIGHT: In this paper, we propose a graph neural network to detect objects from a LiDAR point cloud.

171, TITLE: Self-Learning Video Rain Streak Removal: When Cyclic Consistency Meets Temporal Correspondence
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Self-
Learning_Video_Rain_Streak_Removal_When_Cyclic_Consistency_Meets_Temporal_CVPR_2020_paper.html
AUTHORS: Wenhan Yang, Robby T. Tan, Shiqi Wang, Jiaying Liu
HIGHLIGHT: In this paper, we address the problem of rain streaks removal in video by developing a self-learned rain streak
removal method, which does not require any clean groundtruth images in the training process.

172, TITLE: Neuromorphic Camera Guided High Dynamic Range Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_Neuromorphic_Camera_Guided_High_Dynamic_Range_Imaging_CVP
R_2020_paper.html
AUTHORS: Jin Han, Chu Zhou, Peiqi Duan, Yehui Tang, Chang Xu, Chao Xu, Tiejun Huang, Boxin Shi
HIGHLIGHT: In this paper, we propose a neuromorphic camera guided high dynamic range imaging pipeline, and a network
consisting of specially designed modules according to each step in the pipeline, which bridges the domain gaps on resolution, dynamic
range, and color representation between two types of sensors and images.

173, TITLE: Learning in the Frequency Domain
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Learning_in_the_Frequency_Domain_CVPR_2020_paper.html
AUTHORS: Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, Fengbo Ren
HIGHLIGHT: Inspired by digital signal processing theories, we analyze the spectral bias from the frequency perspective and
propose a learning-based frequency selection method to identify the trivial frequency components which can be removed without
accuracy loss.

174, TITLE: Polarized Reflection Removal With Perfect Alignment in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Lei_Polarized_Reflection_Removal_With_Perfect_Alignment_in_the_Wild_
CVPR_2020_paper.html
AUTHORS: Chenyang Lei, Xuhua Huang, Mengdi Zhang, Qiong Yan, Wenxiu Sun, Qifeng Chen
HIGHLIGHT: We present a novel formulation to removing reflection from polarized images in the wild.

175, TITLE: Learning Multiview 3D Point Cloud Registration
http://openaccess.thecvf.com/content_CVPR_2020/html/Gojcic_Learning_Multiview_3D_Point_Cloud_Registration_CVPR_2020_pa
per.html
AUTHORS: Zan Gojcic, Caifa Zhou, Jan D. Wegner, Leonidas J. Guibas, Tolga Birdal
HIGHLIGHT: We present a novel, end-to-end learnable, multiview 3D point cloud registration algorithm.

176, TITLE: A Sparse Resultant Based Method for Efficient Minimal Solvers
http://openaccess.thecvf.com/content_CVPR_2020/html/Bhayani_A_Sparse_Resultant_Based_Method_for_Efficient_Minimal_Solve
rs_CVPR_2020_paper.html
AUTHORS: Snehal Bhayani, Zuzana Kukelova, Janne Heikkila
HIGHLIGHT: In this paper we study an alternative algebraic method for solving systems of polynomial equations, i.e., the
sparse resultant-based method and propose a novel approach to convert the resultant constraint to an eigenvalue problem.

177, TITLE: Zero-Reference Deep Curve Estimation for Low-Light Image Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Zero-Reference_Deep_Curve_Estimation_for_Low-
Light_Image_Enhancement_CVPR_2020_paper.html
AUTHORS: Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, Runmin Cong
HIGHLIGHT: The paper presents a novel method, Zero-Reference Deep Curve Estimation (Zero-DCE), which formulates
light enhancement as a task of image-specific curve estimation with a deep network.

178, TITLE: BlendedMVS: A Large-Scale Dataset for Generalized Multi-View Stereo Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Yao_BlendedMVS_A_Large-Scale_Dataset_for_Generalized_Multi-
View_Stereo_Networks_CVPR_2020_paper.html
AUTHORS: Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, Long Quan
HIGHLIGHT: In this paper, we introduce BlendedMVS, a novel large-scale dataset, to provide sufficient training ground truth
for learning-based MVS.

179, TITLE: Convolution in the Cloud: Learning Deformable Kernels in 3D Graph Convolution Networks for Point Cloud Analysis
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Convolution_in_the_Cloud_Learning_Deformable_Kernels_in_3D_Graph_CVPR_2020_paper.html
AUTHORS: Zhi-Hao Lin, Sheng-Yu Huang, Yu-Chiang Frank Wang
HIGHLIGHT: In this paper, we propose 3D Graph Convolution Networks (3D-GCN), which is designed to extract local 3D features from point clouds across scales, while shift and scale-invariance properties are introduced.

180, TITLE: A Semi-Supervised Assessor of Neural Architectures
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_A_Semi-Supervised_Assessor_of_Neural_Architectures_CVPR_2020_paper.html
AUTHORS: Yehui Tang, Yunhe Wang, Yixing Xu, Hanting Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, Chang Xu
HIGHLIGHT: In contrast with classical performance predictor optimized in a fully supervised way, this paper suggests a semi-supervised assessor of neural architectures.

181, TITLE: Learning a Reinforced Agent for Flexible Exposure Bracketing Selection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Learning_a_Reinforced_Agent_for_Flexible_Exposure_Bracketing_Selection_CVPR_2020_paper.html
AUTHORS: Zhouxia Wang, Jiawei Zhang, Mude Lin, Jiong Wang, Ping Luo, Jimmy Ren
HIGHLIGHT: Unlike previous methods that have many restrictions such as requiring camera response function, sensor noise model, and a stream of preview images with different exposures (not accessible in some scenarios e.g. mobile applications), we propose a novel deep neural network to automatically select exposure bracketing, named EBSNet, which is sufficiently flexible without having the above restrictions.

182, TITLE: CARS: Continuous Evolution for Efficient Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_CARS_Continuous_Evolution_for_Efficient_Neural_Architecture_Search_CVPR_2020_paper.html
AUTHORS: Zhaohui Yang, Yunhe Wang, Xinghao Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, Chang Xu
HIGHLIGHT: In contrast, we develop an efficient continuous evolutionary approach for searching neural networks.

183, TITLE: Joint 3D Instance Segmentation and Object Detection for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Joint_3D_Instance_Segmentation_and_Object_Detection_for_Autonomous_Driving_CVPR_2020_paper.html
AUTHORS: Dingfu Zhou, Jin Fang, Xibin Song, Liu Liu, Junbo Yin, Yuchao Dai, Hongdong Li, Ruigang Yang
HIGHLIGHT: To tackle this problem, we propose a simple but practical detection framework to jointly predict the 3D BBox and instance segmentation.

184, TITLE: View-GCN: View-Based Graph Convolutional Network for 3D Shape Analysis
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_View-GCN_View-Based_Graph_Convolutional_Network_for_3D_Shape_Analysis_CVPR_2020_paper.html
AUTHORS: Xin Wei, Ruixuan Yu, Jian Sun
HIGHLIGHT: In this work, we propose a novel view-based Graph Convolutional Neural Network, dubbed as view-GCN, to recognize 3D shape based on graph representation of multiple views in flexible view configurations.

185, TITLE: Collaborative Distillation for Ultra-Resolution Universal Style Transfer
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Collaborative_Distillation_for_Ultra-Resolution_Universal_Style_Transfer_CVPR_2020_paper.html
AUTHORS: Huan Wang, Yijun Li, Yuehai Wang, Haoji Hu, Ming-Hsuan Yang
HIGHLIGHT: In this work, we present a new knowledge distillation method (named Collaborative Distillation) for encoder-decoder based neural style transfer to reduce the convolutional filters.

186, TITLE: TomoFluid: Reconstructing Dynamic Fluid From Sparse View Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Zang_TomoFluid_Reconstructing_Dynamic_Fluid_From_Sparse_View_Videos_CVPR_2020_paper.html
AUTHORS: Guangming Zang, Ramzi Idoughi, Congli Wang, Anthony Bennett, Jianguo Du, Scott Skeen, William L. Roberts, Peter Wonka, Wolfgang Heidrich
HIGHLIGHT: In this paper, we present a state-of-the-art 4D tomographic reconstruction framework that integrates several regularizers into a multi-scale matrix free optimization algorithm.

187, TITLE: Instance Shadow Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Instance_Shadow_Detection_CVPR_2020_paper.html
AUTHORS: Tianyu Wang, Xiaowei Hu, Qiong Wang, Pheng-Ann Heng, Chi-Wing Fu

21

HIGHLIGHT: Second, we design LISA, named after Light-guided Instance Shadow-object Association, an end-to-end framework to automatically predict the shadow and object instances, together with the shadow-object associations and light direction. Then, we pair up the predicted shadow and object instances, and match them with the predicted shadow-object associations to generate the final results.

188, TITLE: Self2Self With Dropout: Learning Self-Supervised Denoising From Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Quan_Self2Self_With_Dropout_Learning_Self-Supervised_Denoising_From_Single_Image_CVPR_2020_paper.html
AUTHORS: Yuhui Quan, Mingqin Chen, Tongyao Pang, Hui Ji
HIGHLIGHT: Taking one step further, this paper proposes a self-supervised learning method which only uses the input noisy image itself for training.

189, TITLE: Discrete Model Compression With Resource Constraint for Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Discrete_Model_Compression_With_Resource_Constraint_for_Deep_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Shangqian Gao, Feihu Huang, Jian Pei, Heng Huang
HIGHLIGHT: In this paper, we target to address the problem of compression and acceleration of Convolutional Neural Networks (CNNs).

190, TITLE: Structured Compression by Weight Encryption for Unstructured Pruning and Quantization
http://openaccess.thecvf.com/content_CVPR_2020/html/Kwon_Structured_Compression_by_Weight_Encryption_for_Unstructured_Pruning_and_Quantization_CVPR_2020_paper.html
AUTHORS: Se Jung Kwon, Dongsoo Lee, Byeongwook Kim, Parichay Kapoor, Baeseong Park, Gu-Yeon Wei
HIGHLIGHT: This paper proposes a new weight representation scheme for Sparse Quantized Neural Networks, specifically achieved by fine-grained and unstructured pruning method.

191, TITLE: End-to-End Learning Local Multi-View Descriptors for 3D Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_End-to-End_Learning_Local_Multi-View_Descriptors_for_3D_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Lei Li, Siyu Zhu, Hongbo Fu, Ping Tan, Chiew-Lan Tai
HIGHLIGHT: In this work, we propose an end-to-end framework to learn local multi-view descriptors for 3D point clouds.

192, TITLE: Minimal Solutions for Relative Pose With a Single Affine Correspondence
http://openaccess.thecvf.com/content_CVPR_2020/html/Guan_Minimal_Solutions_for_Relative_Pose_With_a_Single_Affine_Correspondence_CVPR_2020_paper.html
AUTHORS: Banglei Guan, Ji Zhao, Zhang Li, Fang Sun, Friedrich Fraundorfer
HIGHLIGHT: In this paper we present four cases of minimal solutions for two-view relative pose estimation by exploiting the affine transformation between feature points and we demonstrate efficient solvers for these cases.

193, TITLE: Point Cloud Completion by Skip-Attention Network With Hierarchical Folding
http://openaccess.thecvf.com/content_CVPR_2020/html/Wen_Point_Cloud_Completion_by_Skip-Attention_Network_With_Hierarchical_Folding_CVPR_2020_paper.html
AUTHORS: Xin Wen, Tianyang Li, Zhizhong Han, Yu-Shen Liu
HIGHLIGHT: To address this problem, we propose Skip-Attention Network (SA-Net) for 3D point cloud completion.

194, TITLE: Fast-MVSNet: Sparse-to-Dense Multi-View Stereo With Learned Propagation and Gauss-Newton Refinement
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Fast-MVSNet_Sparse-to-Dense_Multi-View_Stereo_With_Learned_Propagation_and_Gauss-Newton_Refinement_CVPR_2020_paper.html
AUTHORS: Zehao Yu, Shenghua Gao
HIGHLIGHT: Towards this end, this paper presents a Fast-MVSNet, a novel sparse-to-dense coarse-to-fine framework, for fast and accurate depth estimation in MVS.

195, TITLE: AANet: Adaptive Aggregation Network for Efficient Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_AANet_Adaptive_Aggregation_Network_for_Efficient_Stereo_Matching_CVPR_2020_paper.html
AUTHORS: Haofei Xu, Juyong Zhang
HIGHLIGHT: In this paper, we aim at completely replacing the commonly used 3D convolutions to achieve fast inference speed while maintaining comparable accuracy.

196, TITLE: Towards Unified INT8 Training for Convolutional Neural Network

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Towards_Unified_INT8_Training_for_Convolutional_Neural_Network_CVPR_2020_paper.html
AUTHORS: Feng Zhu, Ruihao Gong, Fengwei Yu, Xianglong Liu, Yanfei Wang, Zhelong Li, Xiuqi Yang, Junjie Yan
HIGHLIGHT: In this paper, we give an attempt to build a unified 8-bit (INT8) training framework for common convolutional neural networks from the aspects of both accuracy and speed.

197, TITLE: Active 3D Motion Visualization Based on Spatiotemporal Light-Ray Integration
http://openaccess.thecvf.com/content_CVPR_2020/html/Sakaue_Active_3D_Motion_Visualization_Based_on_Spatiotemporal_Light-Ray_Integration_CVPR_2020_paper.html
AUTHORS: Fumihiko Sakaue, Jun Sato
HIGHLIGHT: In this paper, we propose a method of visualizing 3D motion with zero latency.

198, TITLE: Block-Wisely Supervised Neural Architecture Search With Knowledge Distillation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Block-Wisely_Supervised_Neural_Architecture_Search_With_Knowledge_Distillation_CVPR_2020_paper.html
AUTHORS: Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, Xiaojun Chang
HIGHLIGHT: In this work, we propose to modularize the large search space of NAS into blocks to ensure that the potential candidate architectures are fully trained; this reduces the representation shift caused by the shared parameters and leads to the correct rating of the candidates.

199, TITLE: GreedyNAS: Towards Fast One-Shot NAS With Greedy Supernet
http://openaccess.thecvf.com/content_CVPR_2020/html/You_GreedyNAS_Towards_Fast_One-Shot_NAS_With_Greedy_Supernet_CVPR_2020_paper.html
AUTHORS: Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, Changshui Zhang
HIGHLIGHT: In this paper, instead of covering all paths, we ease the burden of supernet by encouraging it to focus more on evaluation of those potentially-good ones, which are identified using a surrogate portion of validation data.

200, TITLE: Learning Filter Pruning Criteria for Deep Convolutional Neural Networks Acceleration
http://openaccess.thecvf.com/content_CVPR_2020/html/He_Learning_Filter_Pruning_Criteria_for_Deep_Convolutional_Neural_Networks_Acceleration_CVPR_2020_paper.html
AUTHORS: Yang He, Yuhang Ding, Ping Liu, Linchao Zhu, Hanwang Zhang, Yi Yang
HIGHLIGHT: In this paper, we propose Learning Filter Pruning Criteria (LFPC) to solve the above problems.

201, TITLE: DIST: Rendering Deep Implicit Signed Distance Function With Differentiable Sphere Tracing
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_DIST_Rendering_Deep_Implicit_Signed_Distance_Function_With_Differentiable_Sphere_CVPR_2020_paper.html
AUTHORS: Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, Zhaopeng Cui
HIGHLIGHT: We propose a differentiable sphere tracing algorithm to bridge the gap between inverse graphics methods and the recently proposed deep learning based implicit signed distance function.

202, TITLE: Visually Imbalanced Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Visually_Imbalanced_Stereo_Matching_CVPR_2020_paper.html
AUTHORS: Yicun Liu, Jimmy Ren, Jiawei Zhang, Jianbo Liu, Mude Lin
HIGHLIGHT: To avoid such collapse, we propose a solution to recover the stereopsis by a joint guided-view-restoration and stereo-reconstruction framework.

203, TITLE: Mesh-Guided Multi-View Stereo With Pyramid Architecture
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Mesh-Guided_Multi-View_Stereo_With_Pyramid_Architecture_CVPR_2020_paper.html
AUTHORS: Yuesong Wang, Tao Guan, Zhuo Chen, Yawei Luo, Keyang Luo, Lili Ju
HIGHLIGHT: To overcome this difficulty, we propose a mesh-guided MVS method with pyramid architecture, which makes use of the surface mesh obtained from coarse-scale images to guide the reconstruction process.

204, TITLE: BiDet: An Efficient Binarized Object Detector
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_BiDet_An_Efficient_Binarized_Object_Detector_CVPR_2020_paper.html
AUTHORS: Ziwei Wang, Ziyi Wu, Jiwen Lu, Jie Zhou
HIGHLIGHT: In this paper, we propose a binarized neural network learning method called BiDet for efficient object detection.

205, TITLE: Local Non-Rigid Structure-From-Motion From Diffeomorphic Mappings

http://openaccess.thecvf.com/content_CVPR_2020/html/Parashar_Local_Non-Rigid_Structure-From-Motion_From_Diffeomorphic_Mappings_CVPR_2020_paper.html
AUTHORS: Shaifali Parashar, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: We propose a new formulation to non-rigid structure-from-motion that only requires the deforming surface to preserve its differential structure.

206, TITLE: Seeing Around Street Corners: Non-Line-of-Sight Detection and Tracking In-the-Wild Using Doppler Radar
http://openaccess.thecvf.com/content_CVPR_2020/html/Scheiner_Seeing_Around_Street_Corners_Non-Line-of-Sight_Detection_and_Tracking_In-the-Wild_Using_CVPR_2020_paper.html
AUTHORS: Nicolas Scheiner, Florian Kraus, Fangyin Wei, Buu Phan, Fahim Mannan, Nils Appenrodt, Werner Ritter, Jurgen Dickmann, Klaus Dietmayer, Bernhard Sick, Felix Heide
HIGHLIGHT: In this work, we depart from visible-wavelength approaches and demonstrate detection, classification, and tracking of hidden objects in large-scale dynamic environments using Doppler radars that can be manufactured at low-cost in series production.

207, TITLE: APQ: Joint Search for Network Architecture, Pruning and Quantization Policy
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_APQ_Joint_Search_for_Network_Architecture_Pruning_and_Quantization_Policy_CVPR_2020_paper.html
AUTHORS: Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, Song Han
HIGHLIGHT: We present APQ, a novel design methodology for efficient deep learning deployment.

208, TITLE: On the Acceleration of Deep Learning Model Parallelism With Staleness
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_On_the_Acceleration_of_Deep_Learning_Model_Parallelism_With_Staleness_CVPR_2020_paper.html
AUTHORS: An Xu, Zhouyuan Huo, Heng Huang
HIGHLIGHT: In this paper, we propose Layer-wise Staleness and a novel efficient training algorithm, Diversely Stale Parameters (DSP), to address these challenges.

209, TITLE: RevealNet: Seeing Behind Objects in RGB-D Scans
http://openaccess.thecvf.com/content_CVPR_2020/html/Hou_RevealNet_Seeing_Behind_Objects_in_RGB-D_Scans_CVPR_2020_paper.html
AUTHORS: Ji Hou, Angela Dai, Matthias Niessner
HIGHLIGHT: We tackle this problem by introducing RevealNet, a new data-driven approach that jointly detects object instances and predicts their complete geometry.

210, TITLE: MemNAS: Memory-Efficient Neural Architecture Search With Grow-Trim Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_MemNAS_Memory-Efficient_Neural_Architecture_Search_With_Grow-Trim_Learning_CVPR_2020_paper.html
AUTHORS: Peiye Liu, Bo Wu, Huadong Ma, Mingoo Seok
HIGHLIGHT: To address this challenge, we propose MemNAS, a novel growing and trimming based neural architecture search framework that optimizes not only performance but also memory requirement of an inference network.

211, TITLE: StegaStamp: Invisible Hyperlinks in Physical Photographs
http://openaccess.thecvf.com/content_CVPR_2020/html/Tancik_StegaStamp_Invisible_Hyperlinks_in_Physical_Photographs_CVPR_2020_paper.html
AUTHORS: Matthew Tancik, Ben Mildenhall, Ren Ng
HIGHLIGHT: Our key technical contribution is StegaStamp, a learned steganographic algorithm to enable robust encoding and decoding of arbitrary hyperlink bitstrings into photos in a manner that approaches perceptual invisibility.

212, TITLE: L2-GCN: Layer-Wise and Learned Efficient Training of Graph Convolutional Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/You_L2-GCN_Layer-Wise_and_Learned_Efficient_Training_of_Graph_Convolutional_Networks_CVPR_2020_paper.html
AUTHORS: Yuning You, Tianlong Chen, Zhangyang Wang, Yang Shen
HIGHLIGHT: In this paper, we propose a novel efficient layer-wise training framework for GCN (L-GCN), that disentangles feature aggregation and feature transformation during training, hence greatly reducing time and memory complexities.

213, TITLE: Polarized Non-Line-of-Sight Imaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Tanaka_Polarized_Non-Line-of-Sight_Imaging_CVPR_2020_paper.html
AUTHORS: Kenichiro Tanaka, Yasuhiro Mukaigawa, Achuta Kadambi
HIGHLIGHT: This paper presents a method of passive non-line-of-sight (NLOS) imaging using polarization cues.

214, TITLE:		AdaBits: Neural Network Quantization With Adaptive Bit-Widths
http://openaccess.thecvf.com/content_CVPR_2020/html/Jin_AdaBits_Neural_Network_Quantization_With_Adaptive_Bit-Widths_CVPR_2020_paper.html
AUTHORS:		Qing Jin, Linjie Yang, Zhenyu Liao
HIGHLIGHT:		In this paper, we investigate a novel option to achieve this goal by enabling adaptive bit-widths of weights and activations in the model.

215, TITLE:		Multi-Scale Boosted Dehazing Network With Dense Feature Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Multi-Scale_Boosted_Dehazing_Network_With_Dense_Feature_Fusion_CVPR_2020_paper.html
AUTHORS:		Hang Dong, Jinshan Pan, Lei Xiang, Zhe Hu, Xinyi Zhang, Fei Wang, Ming-Hsuan Yang
HIGHLIGHT:		In this paper, we propose a Multi-Scale Boosted Dehazing Network with Dense Feature Fusion based on the U-Net architecture.

216, TITLE:		ClusterVO: Clustering Moving Instances and Estimating Visual Odometry for Self and Surroundings
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_ClusterVO_Clustering_Moving_Instances_and_Estimating_Visual_Odometry_for_Self_CVPR_2020_paper.html
AUTHORS:		Jiahui Huang, Sheng Yang, Tai-Jiang Mu, Shi-Min Hu
HIGHLIGHT:		We present ClusterVO, a stereo Visual Odometry which simultaneously clusters and estimates the motion of both ego and surrounding rigid clusters/objects.

217, TITLE:		Automatic Neural Network Compression by Sparsity-Quantization Joint Learning: A Constrained Optimization-Based Approach
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Automatic_Neural_Network_Compression_by_Sparsity-Quantization_Joint_Learning_A_Constrained_CVPR_2020_paper.html
AUTHORS:		Haichuan Yang, Shupeng Gui, Yuhao Zhu, Ji Liu
HIGHLIGHT:		In this paper, we propose a framework to jointly prune and quantize the DNNs automatically according to a target model size without using any hyper-parameters to manually set the compression ratio for each layer.

218, TITLE:		Normal Assisted Stereo Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kusupati_Normal_Assisted_Stereo_Depth_Estimation_CVPR_2020_paper.html
AUTHORS:		Uday Kusupati, Shuo Cheng, Rui Chen, Hao Su
HIGHLIGHT:		In this paper, we study how to enforce the consistency between surface normal and depth at training time to improve the performance.

219, TITLE:		Fusing Wearable IMUs With Multi-View Images for Human Pose Estimation: A Geometric Approach
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Fusing_Wearable_IMUs_With_Multi-View_Images_for_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS:		Zhe Zhang, Chunyu Wang, Wenhu Qin, Wenjun Zeng
HIGHLIGHT:		We present a geometric approach to reinforce the visual features of each pair of joints based on the IMUs.

220, TITLE:		gDLS*: Generalized Pose-and-Scale Estimation Given Scale and Gravity Priors
http://openaccess.thecvf.com/content_CVPR_2020/html/Fragoso_gDLS_Generalized_Pose-and-Scale_Estimation_Given_Scale_and_Gravity_Priors_CVPR_2020_paper.html
AUTHORS:		Victor Fragoso, Joseph DeGol, Gang Hua
HIGHLIGHT:		We present gDLS*, a generalized-camera-model pose-and-scale estimator that utilizes rotation and scale priors.

221, TITLE:		Embodied Language Grounding With 3D Visual Feature Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Prabhudesai_Embodied_Language_Grounding_With_3D_Visual_Feature_Representations_CVPR_2020_paper.html
AUTHORS:		Mihir Prabhudesai, Hsiao-Yu Fish Tung, Syed Ashar Javed, Maximilian Sieb, Adam W. Harley, Katerina Fragkiadaki
HIGHLIGHT:		We present generative models that condition on the dependency tree of an utterance and generate a corresponding visual 3D feature map as well as reason about its plausibility, and detector models that condition on both the dependency tree of an utterance and a related image and localize the object referents in the 3D feature map inferred from the image.

222, TITLE:		Learning to Autofocus
http://openaccess.thecvf.com/content_CVPR_2020/html/Herrmann_Learning_to_Autofocus_CVPR_2020_paper.html
AUTHORS:		Charles Herrmann, Richard Strong Bowen, Neal Wadhwa, Rahul Garg, Qiurui He, Jonathan T. Barron, Ramin Zabih

HIGHLIGHT: We propose a learning-based approach to this problem, and provide a realistic dataset of sufficient size for effective learning.

223, TITLE: Joint Demosaicing and Denoising With Self Guidance
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Joint_Demosaicing_and_Denoising_With_Self_Guidance_CVPR_2020_paper.html
AUTHORS: Lin Liu, Xu Jia, Jianzhuang Liu, Qi Tian
HIGHLIGHT: In this paper, we propose a self-guidance network (SGNet), where the green channels are initially estimated and then works as a guidance to recover all missing values in the input image.

224, TITLE: Forward and Backward Information Retention for Accurate Binary Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Qin_Forward_and_Backward_Information_Retention_for_Accurate_Binary_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, Jingkuan Song
HIGHLIGHT: To address these issues, we propose an Information Retention Network (IR-Net) to retain the information that consists in the forward activations and backward gradients.

225, TITLE: Light Field Spatial Super-Resolution via Deep Combinatorial Geometry Embedding and Structural Consistency Regularization
http://openaccess.thecvf.com/content_CVPR_2020/html/Jin_Light_Field_Spatial_Super-Resolution_via_Deep_Combinatorial_Geometry_Embedding_and_CVPR_2020_paper.html
AUTHORS: Jing Jin, Junhui Hou, Jie Chen, Sam Kwong
HIGHLIGHT: In this paper, we propose a novel learning-based LF spatial SR framework, in which each view of an LF image is first individually super-resolved by exploring the complementary information among views with combinatorial geometry embedding.

226, TITLE: A Multi-Hypothesis Approach to Color Constancy
http://openaccess.thecvf.com/content_CVPR_2020/html/Hernandez-Juarez_A_Multi-Hypothesis_Approach_to_Color_Constancy_CVPR_2020_paper.html
AUTHORS: Daniel Hernandez-Juarez, Sarah Parisot, Benjamin Busam, Ales Leonardis, Gregory Slabaugh, Steven McDonagh
HIGHLIGHT: We propose a Bayesian framework that naturally handles color constancy ambiguity via a multi-hypothesis strategy.

227, TITLE: Learning to Restore Low-Light Images via Decomposition-and-Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Learning_to_Restore_Low-Light_Images_via_Decomposition-and-Enhancement_CVPR_2020_paper.html
AUTHORS: Ke Xu, Xin Yang, Baocai Yin, Rynson W.H. Lau
HIGHLIGHT: Based on this model, we present a novel network that first learns to recover image objects in the low-frequency layer and then enhances high-frequency details based on the recovered image objects.

228, TITLE: Background Matting: The World Is Your Green Screen
http://openaccess.thecvf.com/content_CVPR_2020/html/Sengupta_Background_Matting_The_World_Is_Your_Green_Screen_CVPR_2020_paper.html
AUTHORS: Soumyadip Sengupta, Vivek Jayaram, Brian Curless, Steven M. Seitz, Ira Kemelmacher-Shlizerman
HIGHLIGHT: We propose a method for creating a matte - the per-pixel foreground color and alpha - of a person by taking photos or videos in an everyday setting with a handheld camera.

229, TITLE: Supervised Raw Video Denoising With a Benchmark Dataset on Dynamic Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Yue_Supervised_Raw_Video_Denoising_With_a_Benchmark_Dataset_on_Dynamic_CVPR_2020_paper.html
AUTHORS: Huanjing Yue, Cong Cao, Lei Liao, Ronghe Chu, Jingyu Yang
HIGHLIGHT: In this paper, we solve this problem by creating motions for controllable objects, such as toys, and capturing each static moment for multiple times to generate clean video frames.

230, TITLE: Photometric Stereo via Discrete Hypothesis-and-Test Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Enomoto_Photometric_Stereo_via_Discrete_Hypothesis-and-Test_Search_CVPR_2020_paper.html
AUTHORS: Kenji Enomoto, Michael Waechter, Kiriakos N. Kutulakos, Yasuyuki Matsushita
HIGHLIGHT: In this paper, we consider the problem of estimating surface normals of a scene with spatially varying, general BRDFs observed by a static camera under varying, known, distant illumination.

231, TITLE: Dynamic Convolutions: Exploiting Spatial Sparsity for Faster Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Verelst_Dynamic_Convolutions_Exploiting_Spatial_Sparsity_for_Faster_Inference_CVPR_2020_paper.html
AUTHORS: Thomas Verelst, Tinne Tuytelaars
HIGHLIGHT: To address this inefficiency, we propose a method to dynamically apply convolutions conditioned on the input image.

232, TITLE: Fixed-Point Back-Propagation Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Fixed-Point_Back-Propagation_Training_CVPR_2020_paper.html
AUTHORS: Xishan Zhang, Shaoli Liu, Rui Zhang, Chang Liu, Di Huang, Shiyi Zhou, Jiaming Guo, Qi Guo, Zidong Du, Tian Zhi, Yunji Chen
HIGHLIGHT: In this paper, we propose a novel training approach, which applies a layer-wise precision-adaptive quantization in deep neural networks.

233, TITLE: Heterogeneous Knowledge Distillation Using Information Flow Modeling
http://openaccess.thecvf.com/content_CVPR_2020/html/Passalis_Heterogeneous_Knowledge_Distillation_Using_Information_Flow_Modeling_CVPR_2020_paper.html
AUTHORS: Nikolaos Passalis, Maria Tzelepi, Anastasios Tefas
HIGHLIGHT: In this paper we propose a novel KD method that works by modeling the information flow through the various layers of the teacher model and then train a student model to mimic this information flow.

234, TITLE: Rethinking Differentiable Search for Mixed-Precision Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Cai_Rethinking_Differentiable_Search_for_Mixed-Precision_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Zhaowei Cai, Nuno Vasconcelos
HIGHLIGHT: In this work, the problem of optimal mixed-precision network search (MPS) is considered.

235, TITLE: Residual Feature Aggregation Network for Image Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Residual_Feature_Aggregation_Network_for_Image_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Jie Liu, Wenjie Zhang, Yuting Tang, Jie Tang, Gangshan Wu
HIGHLIGHT: To address this issue, we propose a novel residual feature aggregation (RFA) framework for more efficient feature extraction.

236, TITLE: Resolution Adaptive Networks for Efficient Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Resolution_Adaptive_Networks_for_Efficient_Inference_CVPR_2020_paper.html
AUTHORS: Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, Gao Huang
HIGHLIGHT: In this paper, we focus on spatial redundancy of input samples and propose a novel Resolution Adaptive Network (RANet), which is inspired by the intuition that low-resolution representations are sufficient for classifying "easy" inputs containing large objects with prototypical features, while only some "hard" samples need spatially detailed information.

237, TITLE: Learning to Forget for Meta-Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Baik_Learning_to_Forget_for_Meta-Learning_CVPR_2020_paper.html
AUTHORS: Sungyong Baik, Seokil Hong, Kyoung Mu Lee
HIGHLIGHT: Thus, we propose task-and-layer-wise attenuation on the compromised initialization to reduce its influence.

238, TITLE: Deep Learning for Handling Kernel/model Uncertainty in Image Deconvolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Nan_Deep_Learning_for_Handling_Kernelmodel_Uncertainty_in_Image_Deconvolution_CVPR_2020_paper.html
AUTHORS: Yuesong Nan, Hui Ji
HIGHLIGHT: Based on an error-in-variable (EIV) model of image blurring that takes kernel error into consideration, this paper presents a deep learning method for deconvolution, which unrolls a total-least-squares (TLS) estimator whose relating priors are learned by neural networks (NNs).

239, TITLE: Reflection Scene Separation From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Wan_Reflection_Scene_Separation_From_a_Single_Image_CVPR_2020_paper.html
AUTHORS: Renjie Wan, Boxin Shi, Haoliang Li, Ling-Yu Duan, Alex C. Kot

HIGHLIGHT: In this paper, instead of removing reflection components from the mixture image, we aim at recovering reflection scenes from the mixture image.

240, TITLE: Wavelet Synthesis Net for Disparity Estimation to Synthesize DSLR Calibre Bokeh Effect on Smartphones
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Wavelet_Synthesis_Net_for_Disparity_Estimation_to_Synthesize_DSLR_Calibre_CVPR_2020_paper.html
AUTHORS: Chenchi Luo, Yingmao Li, Kaimo Lin, George Chen, Seok-Jun Lee, Jihwan Choi, Youngjun Francis Yoo, Michael O. Polley
HIGHLIGHT: Empowered by a novel wavelet synthesis network architecture, we have greatly narrowed the gap between DSLR and smartphone camera in terms of the bokeh more than ever before.

241, TITLE: Bundle Adjustment on a Graph Processor
http://openaccess.thecvf.com/content_CVPR_2020/html/Ortiz_Bundle_Adjustment_on_a_Graph_Processor_CVPR_2020_paper.html
AUTHORS: Joseph Ortiz, Mark Pupilli, Stefan Leutenegger, Andrew J. Davison
HIGHLIGHT: We show for the first time that the classical computer vision problem of bundle adjustment (BA) can be solved extremely fast on a graph processor using Gaussian Belief Propagation.

242, TITLE: 3D-ZeF: A 3D Zebrafish Tracking Benchmark Dataset
http://openaccess.thecvf.com/content_CVPR_2020/html/Pedersen_3D-ZeF_A_3D_Zebrafish_Tracking_Benchmark_Dataset_CVPR_2020_paper.html
AUTHORS: Malte Pedersen, Joakim Bruslund Haurum, Stefan Hein Bengtson, Thomas B. Moeslund
HIGHLIGHT: In this work we present a novel publicly available stereo based 3D RGB dataset for multi-object zebrafish tracking, called 3D-ZeF.

243, TITLE: PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Menon_PULSE_Self-Supervised_Photo_Upsampling_via_Latent_Space_Exploration_of_Generative_CVPR_2020_paper.html
AUTHORS: Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, Cynthia Rudin
HIGHLIGHT: We present a novel super-resolution algorithm addressing this problem, PULSE (Photo Upsampling via Latent Space Exploration), which generates high-resolution, realistic images at resolutions previously unseen in the literature.

244, TITLE: Scalability in Perception for Autonomous Driving: Waymo Open Dataset
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Scalability_in_Perception_for_Autonomous_Driving_Waymo_Open_Dataset_CVPR_2020_paper.html
AUTHORS: Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, Dragomir Anguelov
HIGHLIGHT: In an effort to help align the research community's contributions with real-world self-driving problems, we introduce a new large scale, high quality, diverse dataset.

245, TITLE: Extreme Relative Pose Network Under Hybrid Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Extreme_Relative_Pose_Network_Under_Hybrid_Representations_CVPR_2020_paper.html
AUTHORS: Zhenpei Yang, Siming Yan, Qixing Huang
HIGHLIGHT: In this paper, we introduce a novel RGB-D based relative pose estimation approach that is suitable for small-overlapping or non-overlapping scans and can output multiple relative poses.

246, TITLE: Single-Shot Monocular RGB-D Imaging Using Uneven Double Refraction
http://openaccess.thecvf.com/content_CVPR_2020/html/Meuleman_Single-Shot_Monocular_RGB-D_Imaging_Using_Uneven_Double_Refraction_CVPR_2020_paper.html
AUTHORS: Andreas Meuleman, Seung-Hwan Baek, Felix Heide, Min H. Kim
HIGHLIGHT: In this work, we propose a method for monocular single-shot RGB-D imaging.

247, TITLE: Inverse Rendering for Complex Indoor Scenes: Shape, Spatially-Varying Lighting and SVBRDF From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Inverse_Rendering_for_Complex_Indoor_Scenes_Shape_Spatially-Varying_Lighting_and_CVPR_2020_paper.html
AUTHORS: Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, Manmohan Chandraker
HIGHLIGHT: We propose a deep inverse rendering framework for indoor scenes.

248, TITLE: 3D Packing for Self-Supervised Monocular Depth Estimation

http://openaccess.thecvf.com/content_CVPR_2020/html/Guizilini_3D_Packing_for_Self-Supervised_Monocular_Depth_Estimation_CVPR_2020_paper.html
AUTHORS: Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, Adrien Gaidon
HIGHLIGHT: In this work, we propose a novel self-supervised monocular depth estimation method combining geometry with a new deep network, PackNet, learned only from unlabeled monocular videos.

249, TITLE: Cascade Cost Volume for High-Resolution Multi-View Stereo and Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Gu_Cascade_Cost_Volume_for_High-Resolution_Multi-View_Stereo_and_Stereo_Matching_CVPR_2020_paper.html
AUTHORS: Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, Ping Tan
HIGHLIGHT: In this paper, we propose a both memory and time efficient cost volume formulation that is complementary to existing multi-view stereo and stereo matching approaches based on 3D cost volumes.

250, TITLE: From Two Rolling Shutters to One Global Shutter
http://openaccess.thecvf.com/content_CVPR_2020/html/Albl_From_Two_Rolling_Shutters_to_One_Global_Shutter_CVPR_2020_paper.html
AUTHORS: Cenek Albl, Zuzana Kukelova, Viktor Larsson, Michal Polic, Tomas Pajdla, Konrad Schindler
HIGHLIGHT: We explore a surprisingly simple camera configuration that makes it possible to undo the rolling shutter distortion: two cameras mounted to have different rolling shutter directions.

251, TITLE: Deep Global Registration
http://openaccess.thecvf.com/content_CVPR_2020/html/Choy_Deep_Global_Registration_CVPR_2020_paper.html
AUTHORS: Christopher Choy, Wei Dong, Vladlen Koltun
HIGHLIGHT: We present Deep Global Registration, a differentiable framework for pairwise registration of real-world 3D scans.

252, TITLE: Deep Stereo Using Adaptive Thin Volume Representation With Uncertainty Awareness
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Deep_Stereo_Using_Adaptive_Thin_Volume_Representation_With_Uncertainty_Awareness_CVPR_2020_paper.html
AUTHORS: Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, Hao Su
HIGHLIGHT: We present Uncertainty-aware Cascaded Stereo Network (UCS-Net) for 3D reconstruction from multiple RGB images.

253, TITLE: Why Having 10,000 Parameters in Your Camera Model Is Better Than Twelve
http://openaccess.thecvf.com/content_CVPR_2020/html/Schops_Why_Having_10000_Parameters_in_Your_Camera_Model_Is_Better_CVPR_2020_paper.html
AUTHORS: Thomas Schops, Viktor Larsson, Marc Pollefeys, Torsten Sattler
HIGHLIGHT: We propose a calibration pipeline for generic models that is fully automated, easy to use, and can act as a drop-in replacement for parametric calibration, with a focus on accuracy.

254, TITLE: Blur Aware Calibration of Multi-Focus Plenoptic Camera
http://openaccess.thecvf.com/content_CVPR_2020/html/Labussiere_Blur_Aware_Calibration_of_Multi-Focus_Plenoptic_Camera_CVPR_2020_paper.html
AUTHORS: Mathieu Labussiere, Celine Teuliere, Frederic Bernardin, Omar Ait-Aider
HIGHLIGHT: This paper presents a novel calibration algorithm for Multi-Focus Plenoptic Cameras (MFPCs) using raw images only.

255, TITLE: Learning Fused Pixel and Feature-Based View Reconstructions for Light Fields
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Learning_Fused_Pixel_and_Feature-Based_View_Reconstructions_for_Light_Fields_CVPR_2020_paper.html
AUTHORS: Jinglei Shi, Xiaoran Jiang, Christine Guillemot
HIGHLIGHT: In this paper, we present a learning-based framework for light field view synthesis from a subset of input views.

256, TITLE: SAL: Sign Agnostic Learning of Shapes From Raw Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Atzmon_SAL_Sign_Agnostic_Learning_of_Shapes_From_Raw_Data_CVPR_2020_paper.html
AUTHORS: Matan Atzmon, Yaron Lipman
HIGHLIGHT: In this paper we introduce Sign Agnostic Learning (SAL), a deep learning approach for learning implicit shape representations directly from raw, unsigned geometric data, such as point clouds and triangle soups.

257, TITLE: Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval

http://openaccess.thecvf.com/content_CVPR_2020/html/Weyand_Google_Landmarks_Dataset_v2_-_A_Large-Scale_Benchmark_for_Instance-Level_CVPR_2020_paper.html
AUTHORS: Tobias Weyand, Andre Araujo, Bingyi Cao, Jack Sim
HIGHLIGHT: We introduce the Google Landmarks Dataset v2 (GLDv2), a new benchmark for large-scale, fine-grained instance recognition and image retrieval in the domain of human-made and natural landmarks.

258, TITLE: Instance Guided Proposal Network for Person Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Instance_Guided_Proposal_Network_for_Person_Search_CVPR_2020_paper.html
AUTHORS: Wenkai Dong, Zhaoxiang Zhang, Chunfeng Song, Tieniu Tan
HIGHLIGHT: In this paper, we propose a new detection network for person search, named Instance Guided Proposal Network (IGPN), which can learn the similarity between query persons and proposals.

259, TITLE: Which Is Plagiarism: Fashion Image Retrieval Based on Regional Representation for Design Protection
http://openaccess.thecvf.com/content_CVPR_2020/html/Lang_Which_Is_Plagiarism_Fashion_Image_Retrieval_Based_on_Regional_Representation_CVPR_2020_paper.html
AUTHORS: Yining Lang, Yuan He, Fan Yang, Jianfeng Dong, Hui Xue
HIGHLIGHT: Different from the existing works that mainly focus on identical or similar fashion item retrieval, in this paper, we aim to study the plagiarized clothes retrieval which is somewhat ignored in the academic community while itself has great application value.

260, TITLE: Inter-Task Association Critic for Cross-Resolution Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Inter-Task_Association_Critic_for_Cross-Resolution_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Zhiyi Cheng, Qi Dong, Shaogang Gong, Xiatian Zhu
HIGHLIGHT: In this paper, we introduce a novel model training regularisation method, called Inter-Task Association Critic (INTACT), to address this fundamental problem.

261, TITLE: FineGym: A Hierarchical Video Dataset for Fine-Grained Action Understanding
http://openaccess.thecvf.com/content_CVPR_2020/html/Shao_FineGym_A_Hierarchical_Video_Dataset_for_Fine-Grained_Action_Understanding_CVPR_2020_paper.html
AUTHORS: Dian Shao, Yue Zhao, Bo Dai, Dahua Lin
HIGHLIGHT: To take action recognition to a new level, we develop FineGym, a new dataset built on top of gymnasium videos.

262, TITLE: Mapillary Street-Level Sequences: A Dataset for Lifelong Place Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Warburg_Mapillary_Street-Level_Sequences_A_Dataset_for_Lifelong_Place_Recognition_CVPR_2020_paper.html
AUTHORS: Frederik Warburg, Soren Hauberg, Manuel Lopez-Antequera, Pau Gargallo, Yubin Kuang, Javier Civera
HIGHLIGHT: We contribute with Mapillary Street-Level Sequences (SLS), a large dataset for urban and suburban place recognition from image sequences.

263, TITLE: BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_BDD100K_A_Diverse_Driving_Dataset_for_Heterogeneous_Multitask_Learning_CVPR_2020_paper.html
AUTHORS: Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, Trevor Darrell
HIGHLIGHT: We construct BDD100K, the largest driving video dataset with 100K videos and 10 tasks to evaluate the exciting progress of image recognition algorithms on autonomous driving.

264, TITLE: Rethinking Computer-Aided Tuberculosis Diagnosis
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Rethinking_Computer-Aided_Tuberculosis_Diagnosis_CVPR_2020_paper.html
AUTHORS: Yun Liu, Yu-Huan Wu, Yunfeng Ban, Huifang Wang, Ming-Ming Cheng
HIGHLIGHT: To solve this problem, we establish a large-scale TB dataset, namely Tuberculosis X-ray (TBX11K) dataset.

265, TITLE: IntrA: 3D Intracranial Aneurysm Dataset for Deep Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_IntrA_3D_Intracranial_Aneurysm_Dataset_for_Deep_Learning_CVPR_2020_paper.html
AUTHORS: Xi Yang, Ding Xia, Taichi Kin, Takeo Igarashi
HIGHLIGHT: In this paper, instead of 2D medical images, we introduce an open-access 3D intracranial aneurysm dataset, IntrA, that makes the application of points-based and mesh-based classification and segmentation models available.

266, TITLE: Revisiting Saliency Metrics: Farthest-Neighbor Area Under Curve
http://openaccess.thecvf.com/content_CVPR_2020/html/Jia_Revisiting_Saliency_Metrics_Farthest-Neighbor_Area_Under_Curve_CVPR_2020_paper.html
AUTHORS: Sen Jia, Neil D. B. Bruce
HIGHLIGHT: In this paper, we propose a new metric to address the long-standing problem of center bias in saliency evaluation.

267, TITLE: Computing the Testing Error Without a Testing Set
http://openaccess.thecvf.com/content_CVPR_2020/html/Corneanu_Computing_the_Testing_Error_Without_a_Testing_Set_CVPR_2020_paper.html
AUTHORS: Ciprian A. Corneanu, Sergio Escalera, Aleix M. Martinez
HIGHLIGHT: Here, we derive an algorithm to estimate the performance gap between training and testing without the need of a testing dataset.

268, TITLE: Improving Confidence Estimates for Unfamiliar Examples
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Improving_Confidence_Estimates_for_Unfamiliar_Examples_CVPR_2020_paper.html
AUTHORS: Zhizhong Li, Derek Hoiem
HIGHLIGHT: In this paper, we compare and evaluate several methods to improve confidence estimates for unfamiliar and familiar samples.

269, TITLE: CycleISP: Real Image Restoration via Improved Data Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Zamir_CycleISP_Real_Image_Restoration_via_Improved_Data_Synthesis_CVPR_2020_paper.html
AUTHORS: Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, Ling Shao
HIGHLIGHT: In this paper, we present a framework that models camera imaging pipeline in forward and reverse directions.

270, TITLE: Enhanced Blind Face Restoration With Multi-Exemplar Images and Adaptive Spatial Feature Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Enhanced_Blind_Face_Restoration_With_Multi-Exemplar_Images_and_Adaptive_Spatial_CVPR_2020_paper.html
AUTHORS: Xiaoming Li, Wenyu Li, Dongwei Ren, Hongzhi Zhang, Meng Wang, Wangmeng Zuo
HIGHLIGHT: To address these issues, this paper suggests to enhance blind face restoration performance by utilizing multi-exemplar images and adaptive fusion of features from guidance and degraded images.

271, TITLE: Explorable Super Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Bahat_Explorable_Super_Resolution_CVPR_2020_paper.html
AUTHORS: Yuval Bahat, Tomer Michaeli
HIGHLIGHT: In this paper, we introduce the task of explorable super resolution.

272, TITLE: Syn2Real Transfer Learning for Image Deraining Using Gaussian Processes
http://openaccess.thecvf.com/content_CVPR_2020/html/Yasarla_Syn2Real_Transfer_Learning_for_Image_Deraining_Using_Gaussian_Processes_CVPR_2020_paper.html
AUTHORS: Rajeev Yasarla, Vishwanath A. Sindagi, Vishal M. Patel
HIGHLIGHT: We propose a Gaussian Process-based semi-supervised learning framework which enables the network in learning to derain using synthetic dataset while generalizing better using unlabeled real-world images.

273, TITLE: Deblurring by Realistic Blurring
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Deblurring_by_Realistic_Blurring_CVPR_2020_paper.html
AUTHORS: Kaihao Zhang, Wenhan Luo, Yiran Zhong, Lin Ma, Bjorn Stenger, Wei Liu, Hongdong Li
HIGHLIGHT: To address this problem, we propose a new method which combines two GAN models, i.e., a learning-to-Blur GAN (BGAN) and learning-to-DeBlur GAN (DBGAN), in order to learn a better model for image deblurring by primarily learning how to blur images.

274, TITLE: Bringing Old Photos Back to Life
http://openaccess.thecvf.com/content_CVPR_2020/html/Wan_Bringing_Old_Photos_Back_to_Life_CVPR_2020_paper.html
AUTHORS: Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, Fang Wen
HIGHLIGHT: We propose to restore old photos that suffer from severe degradation through a deep learning approach.

275, TITLE: A Physics-Based Noise Formation Model for Extreme Low-Light Raw Denoising
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_A_Physics-Based_Noise_Formation_Model_for_Extreme_Low-Light_Raw_Denoising_CVPR_2020_paper.html
AUTHORS: Kaixuan Wei, Ying Fu, Jiaolong Yang, Hua Huang
HIGHLIGHT: To address this issue, we present a highly accurate noise formation model based on the characteristics of CMOS photosensors, thereby enabling us to synthesize realistic samples that better match the physics of image formation process.

276, TITLE: Camouflaged Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_Camouflaged_Object_Detection_CVPR_2020_paper.html
AUTHORS: Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, Ling Shao
HIGHLIGHT: We present a comprehensive study on a new task named camouflaged object detection (COD), which aims to identify objects that are "seamlessly" embedded in their surroundings.

277, TITLE: Holistically-Attracted Wireframe Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Xue_Holistically-Attracted_Wireframe_Parsing_CVPR_2020_paper.html
AUTHORS: Nan Xue, Tianfu Wu, Song Bai, Fudong Wang, Gui-Song Xia, Liangpei Zhang, Philip H.S. Torr
HIGHLIGHT: This paper presents a fast and parsimonious parsing method to accurately and robustly detect a vectorized wireframe in an input image with a single forward pass.

278, TITLE: Conv-MPN: Convolutional Message Passing Neural Network for Structured Outdoor Architecture Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Conv-MPN_Convolutional_Message_Passing_Neural_Network_for_Structured_Outdoor_Architecture_CVPR_2020_paper.html
AUTHORS: Fuyang Zhang, Nelson Nauata, Yasutaka Furukawa
HIGHLIGHT: This paper proposes a novel message passing neural (MPN) architecture Conv-MPN, which reconstructs an outdoor building as a planar graph from a single RGB image.

279, TITLE: Domain Adaptation for Image Dehazing
http://openaccess.thecvf.com/content_CVPR_2020/html/Shao_Domain_Adaptation_for_Image_Dehazing_CVPR_2020_paper.html
AUTHORS: Yuanjie Shao, Lerenhan Li, Wenqi Ren, Changxin Gao, Nong Sang
HIGHLIGHT: To address this issue, we propose a domain adaptation paradigm, which consists of an image translation module and two image dehazing modules.

280, TITLE: Auto-Encoding Twin-Bottleneck Hashing
http://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Auto-Encoding_Twin-Bottleneck_Hashing_CVPR_2020_paper.html
AUTHORS: Yuming Shen, Jie Qin, Jiaxin Chen, Mengyang Yu, Li Liu, Fan Zhu, Fumin Shen, Ling Shao
HIGHLIGHT: In this paper, we tackle the above problems by proposing an efficient and adaptive code-driven graph, which is updated by decoding in the context of an auto-encoder.

281, TITLE: Agriculture-Vision: A Large Aerial Image Database for Agricultural Pattern Analysis
http://openaccess.thecvf.com/content_CVPR_2020/html/Chiu_Agriculture-Vision_A_Large_Aerial_Image_Database_for_Agricultural_Pattern_Analysis_CVPR_2020_paper.html
AUTHORS: Mang Tik Chiu, Xingqian Xu, Yunchao Wei, Zilong Huang, Alexander G. Schwing, Robert Brunner, Hrant Khachatrian, Hovnatan Karapetyan, Ivan Dozier, Greg Rose, David Wilson, Adrian Tudor, Naira Hovakimyan, Thomas S. Huang, Honghui Shi
HIGHLIGHT: To encourage research in computer vision for agriculture, we present Agriculture-Vision: a large-scale aerial farmland image dataset for semantic segmentation of agricultural patterns.

282, TITLE: Bi-Directional Interaction Network for Person Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Bi-Directional_Interaction_Network_for_Person_Search_CVPR_2020_paper.html
AUTHORS: Wenkai Dong, Zhaoxiang Zhang, Chunfeng Song, Tieniu Tan
HIGHLIGHT: To address this issue, we propose a Siamese network which owns an additional instance-aware branch, named Bi-directional Interaction Network (BINet).

283, TITLE: Meshlet Priors for 3D Mesh Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Badki_Meshlet_Priors_for_3D_Mesh_Reconstruction_CVPR_2020_paper.html
AUTHORS: Abhishek Badki, Orazio Gallo, Jan Kautz, Pradeep Sen
HIGHLIGHT: We introduce meshlets, small patches of mesh that we use to learn local shape priors.

284, TITLE: Space-Time-Aware Multi-Resolution Video Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Haris_Space-Time-Aware_Multi-Resolution_Video_Enhancement_CVPR_2020_paper.html
AUTHORS: Muhammad Haris, Greg Shakhnarovich, Norimichi Ukita
HIGHLIGHT: We consider the problem of space-time super-resolution (ST-SR): increasing spatial resolution of video frames and simultaneously interpolating frames to increase the frame rate.

285, TITLE: FSS-1000: A 1000-Class Dataset for Few-Shot Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_FSS-1000_A_1000-Class_Dataset_for_Few-Shot_Segmentation_CVPR_2020_paper.html
AUTHORS: Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, Chi-Keung Tang
HIGHLIGHT: In this paper, we are interested in few-shot object segmentation where the number of annotated training examples are limited to 5 only.
To evaluate and validate the performance of our approach, we have built a few-shot segmentation dataset, FSS-1000, which consists of 1000 object classes with pixelwise annotation of ground-truth segmentation.

286, TITLE: MSeg: A Composite Dataset for Multi-Domain Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lambert_MSeg_A_Composite_Dataset_for_Multi-Domain_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: John Lambert, Zhuang Liu, Ozan Sener, James Hays, Vladlen Koltun
HIGHLIGHT: We present MSeg, a composite dataset that unifies se- mantic segmentation datasets from different domains.

287, TITLE: Learning Multi-Granular Hypergraphs for Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Learning_Multi-Granular_Hypergraphs_for_Video-Based_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Yichao Yan, Jie Qin, Jiaxin Chen, Li Liu, Fan Zhu, Ying Tai, Ling Shao
HIGHLIGHT: In this work, we propose a novel graph-based framework, namely Multi-Granular Hypergraph (MGH), to pursue better representational capabilities by modeling spatiotemporal dependencies in terms of multiple granularities.

288, TITLE: Online Joint Multi-Metric Adaptation From Frequent Sharing-Subset Mining for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Online_Joint_Multi-Metric_Adaptation_From_Frequent_Sharing-Subset_Mining_for_Person_CVPR_2020_paper.html
AUTHORS: Jiahuan Zhou, Bing Su, Ying Wu
HIGHLIGHT: Therefore, we propose an online joint multi-metric adaptation model to adapt the offline learned P-RID models for the online data by learning a series of metrics for all the sharing-subsets.

289, TITLE: Taking a Deeper Look at Co-Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_Taking_a_Deeper_Look_at_Co-Salient_Object_Detection_CVPR_2020_paper.html
AUTHORS: Deng-Ping Fan, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Huazhu Fu, Ming-Ming Cheng
HIGHLIGHT: To tackle this issue, we first collect a new high-quality dataset, named CoSOD3k, which contains 3,316 images divided into 160 groups with multiple level annotations, i.e., category, bounding box, object, and instance levels.

290, TITLE: Single-Stage 6D Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Single-Stage_6D_Object_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Yinlin Hu, Pascal Fua, Wei Wang, Mathieu Salzmann
HIGHLIGHT: In this work, we introduce a deep architecture that directly regresses 6D poses from correspondences.

291, TITLE: OccuSeg: Occupancy-Aware 3D Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_OccuSeg_Occupancy-Aware_3D_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Lei Han, Tian Zheng, Lan Xu, Lu Fang
HIGHLIGHT: In this paper, we define "3D occupancy size", as the number of voxels occupied by each instance. It owns advantages of robustness in prediction, on which basis, OccuSeg, an occupancy-aware 3D instance segmentation scheme is proposed.

292, TITLE: Camera Trace Erasing
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Camera_Trace_Erasing_CVPR_2020_paper.html
AUTHORS: Chang Chen, Zhiwei Xiong, Xiaoming Liu, Feng Wu
HIGHLIGHT: In this paper, we address a new low-level vision problem, camera trace erasing, to reveal the weakness of trace-based forensic methods.

293, TITLE:	Deep Metric Learning via Adaptive Learnable Assessment
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Deep_Metric_Learning_via_Adaptive_Learnable_Assessment_CVPR_2020_paper.html
AUTHORS:	Wenzhao Zheng, Jiwen Lu, Jie Zhou
HIGHLIGHT:	In this paper, we propose a deep metric learning via adaptive learnable assessment (DML-ALA) method for image retrieval and clustering, which aims to learn a sample assessment strategy to maximize the generalization of the trained metric.

294, TITLE:	Deep Representation Learning on Long-Tailed Data: A Learnable Embedding Augmentation Perspective
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Deep_Representation_Learning_on_Long-Tailed_Data_A_Learnable_Embedding_Augmentation_CVPR_2020_paper.html
AUTHORS:	Jialun Liu, Yifan Sun, Chuchu Han, Zhaopeng Dou, Wenhui Li
HIGHLIGHT:	To this end, we propose to augment each instance of the tail classes with certain disturbances in the deep feature space.

295, TITLE:	Fantastic Answers and Where to Find Them: Immersive Question-Directed Visual Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Fantastic_Answers_and_Where_to_Find_Them_Immersive_Question-Directed_Visual_CVPR_2020_paper.html
AUTHORS:	Ming Jiang, Shi Chen, Jinhui Yang, Qi Zhao
HIGHLIGHT:	Specifically, we introduce the first dataset of top-down attention in immersive scenes.

296, TITLE:	HUMBI: A Large Multiview Dataset of Human Body Expressions
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_HUMBI_A_Large_Multiview_Dataset_of_Human_Body_Expressions_CVPR_2020_paper.html
AUTHORS:	Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth Venkatesh, Jaesik Park, Jihun Yu, Hyun Soo Park
HIGHLIGHT:	This paper presents a new large multiview dataset called HUMBI for human body expressions with natural clothing.

297, TITLE:	Image Search With Text Feedback by Visiolinguistic Attention Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Image_Search_With_Text_Feedback_by_Visiolinguistic_Attention_Learning_CVPR_2020_paper.html
AUTHORS:	Yanbei Chen, Shaogang Gong, Loris Bazzani
HIGHLIGHT:	In this work, we tackle this task by a novel Visiolinguistic Attention Learning (VAL) framework.

298, TITLE:	Image Processing Using Multi-Code GAN Prior
http://openaccess.thecvf.com/content_CVPR_2020/html/Gu_Image_Processing_Using_Multi-Code_GAN_Prior_CVPR_2020_paper.html
AUTHORS:	Jinjin Gu, Yujun Shen, Bolei Zhou
HIGHLIGHT:	In this work, we propose a novel approach, called mGANprior, to incorporate the well-trained GANs as effective prior to a variety of image processing tasks.

299, TITLE:	What Does Plate Glass Reveal About Camera Calibration?
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_What_Does_Plate_Glass_Reveal_About_Camera_Calibration_CVPR_2020_paper.html
AUTHORS:	Qian Zheng, Jinnan Chen, Zhan Lu, Boxin Shi, Xudong Jiang, Kim-Hui Yap, Ling-Yu Duan, Alex C. Kot
HIGHLIGHT:	This paper aims to calibrate the orientation of glass and the field of view of the camera from a single reflection-contaminated image.
We collect a dataset containing 320 samples as well as their camera parameters for evaluation.

300, TITLE:	Zero-Assignment Constraint for Graph Matching With Outliers
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Zero-Assignment_Constraint_for_Graph_Matching_With_Outliers_CVPR_2020_paper.html
AUTHORS:	Fudong Wang, Nan Xue, Jin-Gang Yu, Gui-Song Xia
HIGHLIGHT:	To address this issue, we present the zero-assignment constraint (ZAC) for approaching the graph matching problem in the presence of outliers.

301, TITLE:	Cascaded Deep Video Deblurring Using Temporal Sharpness Prior
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Cascaded_Deep_Video_Deblurring_Using_Temporal_Sharpness_Prior_CVPR_2020_paper.html
AUTHORS:	Jinshan Pan, Haoran Bai, Jinhui Tang
HIGHLIGHT:	We present a simple and effective deep convolutional neural network (CNN) model for video deblurring.

302, TITLE: JL-DCF: Joint Learning and Densely-Cooperative Fusion Framework for RGB-D Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Fu_JL-DCF_Joint_Learning_and_Densely-Cooperative_Fusion_Framework_for_RGB-D_Salient_CVPR_2020_paper.html
AUTHORS: Keren Fu, Deng-Ping Fan, Ge-Peng Ji, Qijun Zhao
HIGHLIGHT: This paper proposes a novel joint learning and densely-cooperative fusion (JL-DCF) architecture for RGB-D salient object detection.

303, TITLE: From Fidelity to Perceptual Quality: A Semi-Supervised Approach for Low-Light Image Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_From_Fidelity_to_Perceptual_Quality_A_Semi-Supervised_Approach_for_Low-Light_CVPR_2020_paper.html
AUTHORS: Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, Jiaying Liu
HIGHLIGHT: To address these problems, we propose a novel semi-supervised learning approach for low-light image enhancement.

304, TITLE: Unsupervised Adaptation Learning for Hyperspectral Imagery Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Unsupervised_Adaptation_Learning_for_Hyperspectral_Imagery_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Lei Zhang, Jiangtao Nie, Wei Wei, Yanning Zhang, Shengcai Liao, Ling Shao
HIGHLIGHT: To tackle this problem, we present an unsupervised adaptation learning (UAL) framework.

305, TITLE: Central Similarity Quantization for Efficient Image and Video Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Yuan_Central_Similarity_Quantization_for_Efficient_Image_and_Video_Retrieval_CVPR_2020_paper.html
AUTHORS: Li Yuan, Tao Wang, Xiaopeng Zhang, Francis EH Tay, Zequn Jie, Wei Liu, Jiashi Feng
HIGHLIGHT: In this work, we propose a new global similarity metric, termed as central similarity, with which the hash codes of similar data pairs are encouraged to approach a common center and those for dissimilar pairs to converge to different centers, to improve hash learning efficiency and retrieval accuracy.

306, TITLE: ARCH: Animatable Reconstruction of Clothed Humans
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_ARCH_Animatable_Reconstruction_of_Clothed_Humans_CVPR_2020_paper.html
AUTHORS: Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, Tony Tung
HIGHLIGHT: In this paper, we propose ARCH (Animatable Reconstruction of Clothed Humans), a novel end-to-end framework for accurate reconstruction of animation-ready 3D clothed humans from a monocular image.

307, TITLE: A Model-Driven Deep Neural Network for Single Image Rain Removal
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_A_Model-Driven_Deep_Neural_Network_for_Single_Image_Rain_Removal_CVPR_2020_paper.html
AUTHORS: Hong Wang, Qi Xie, Qian Zhao, Deyu Meng
HIGHLIGHT: To this issue, in this paper, we propose a model-driven deep neural network for the task, with fully interpretable network structures.

308, TITLE: Novel Object Viewpoint Estimation Through Reconstruction Alignment
http://openaccess.thecvf.com/content_CVPR_2020/html/Banani_Novel_Object_Viewpoint_Estimation_Through_Reconstruction_Alignment_CVPR_2020_paper.html
AUTHORS: Mohamed El Banani, Jason J. Corso, David F. Fouhey
HIGHLIGHT: The goal of this paper is to estimate the viewpoint for a novel object.

309, TITLE: Creating Something From Nothing: Unsupervised Knowledge Distillation for Cross-Modal Hashing
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Creating_Something_From_Nothing_Unsupervised_Knowledge_Distillation_for_Cross_Modal_Hashing_CVPR_2020_paper.html
AUTHORS: Hengtong Hu, Lingxi Xie, Richang Hong, Qi Tian
HIGHLIGHT: In this paper, we propose a novel approach that enables guiding a supervised method using outputs produced by an unsupervised method.

310, TITLE: Evaluating Weakly Supervised Object Localization Methods Right
http://openaccess.thecvf.com/content_CVPR_2020/html/Choe_Evaluating_Weakly_Supervised_Object_Localization_Methods_Right_CVPR_2020_paper.html
AUTHORS: Junsuk Choe, Seong Joon Oh, Seungho Lee, Sanghyuk Chun, Zeynep Akata, Hyunjung Shim
HIGHLIGHT: In this paper, we argue that WSOL task is ill-posed with only image-level labels, and propose a new evaluation protocol where full supervision is limited to only a small held-out set not overlapping with the test set.

311, TITLE: Style Normalization and Restitution for Generalizable Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Jin_Style_Normalization_and_Restitution_for_Generalizable_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, Li Zhang
HIGHLIGHT: In this paper, we aim to design a generalizable person ReID framework which trains a model on source domains yet is able to generalize/perform well on target domains.

312, TITLE: Reconstruct Locally, Localize Globally: A Model Free Method for Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cai_Reconstruct_Locally_Localize_Globally_A_Model_Free_Method_for_Object_CVPR_2020_paper.html
AUTHORS: Ming Cai, Ian Reid
HIGHLIGHT: Instead, we propose a learning-based method whose input is a collection of images of a target object, and whose output is the pose of the object in a novel view.

313, TITLE: RoboTHOR: An Open Simulation-to-Real Embodied AI Platform
http://openaccess.thecvf.com/content_CVPR_2020/html/Deitke_RoboTHOR_An_Open_Simulation-to-Real_Embodied_AI_Platform_CVPR_2020_paper.html
AUTHORS: Matt Deitke, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Dustin Schwenk, Eli VanderBilt, Matthew Wallingford, Luca Weihs, Mark Yatskar, Ali Farhadi
HIGHLIGHT: In this paper, we introduce RoboTHOR to democratize research in interactive and embodied visual AI.

314, TITLE: All in One Bad Weather Removal Using Architectural Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_All_in_One_Bad_Weather_Removal_Using_Architectural_Search_CVPR_2020_paper.html
AUTHORS: Ruoteng Li, Robby T. Tan, Loong-Fah Cheong
HIGHLIGHT: In this paper, we propose a method that can handle multiple bad weather degradations: rain, fog, snow and adherent raindrops using a single network.

315, TITLE: Relation-Aware Global Attention for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Relation-Aware_Global_Attention_for_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Xin Jin, Zhibo Chen
HIGHLIGHT: In this work, we propose an effective Relation-Aware Global Attention (RGA) module which captures the global structural information for better attention learning.

316, TITLE: HOnnotate: A Method for 3D Annotation of Hand and Object Poses
http://openaccess.thecvf.com/content_CVPR_2020/html/Hampali_HOnnotate_A_Method_for_3D_Annotation_of_Hand_and_Object_CVPR_2020_paper.html
AUTHORS: Shreyas Hampali, Mahdi Rad, Markus Oberweger, Vincent Lepetit
HIGHLIGHT: We propose a method for annotating images of a hand manipulating an object with the 3D poses of both the hand and the object, together with a dataset created using this method.

317, TITLE: Celeb-DF: A Large-Scale Challenging Dataset for DeepFake Forensics
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Celeb-DF_A_Large-Scale_Challenging_Dataset_for_DeepFake_Forensics_CVPR_2020_paper.html
AUTHORS: Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, Siwei Lyu
HIGHLIGHT: We present a new large-scale challenging DeepFake video dataset, Celeb-DF, which contains 5,639 high-quality DeepFake videos of celebrities generated using improved synthesis process.

318, TITLE: Deep Unfolding Network for Image Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Deep_Unfolding_Network_for_Image_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Kai Zhang, Luc Van Gool, Radu Timofte
HIGHLIGHT: To address this issue, this paper proposes an end-to-end trainable unfolding network which leverages both learningbased methods and model-based methods.

319, TITLE: On the Uncertainty of Self-Supervised Monocular Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Poggi_On_the_Uncertainty_of_Self-Supervised_Monocular_Depth_Estimation_CVPR_2020_paper.html
AUTHORS: Matteo Poggi, Filippo Aleotti, Fabio Tosi, Stefano Mattoccia

HIGHLIGHT: Purposely, we explore for the first time how to estimate the uncertainty for this task and how this affects depth accuracy, proposing a novel peculiar technique specifically designed for self-supervised approaches.

320, TITLE: Proxy Anchor Loss for Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Proxy_Anchor_Loss_for_Deep_Metric_Learning_CVPR_2020_paper.html
AUTHORS: Sungyeon Kim, Dongwon Kim, Minsu Cho, Suha Kwak
HIGHLIGHT: This paper presents a new proxy-based loss that takes advantages of both pair- and proxy-based methods and overcomes their limitations.

321, TITLE: Unsupervised Learning for Intrinsic Image Decomposition From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Unsupervised_Learning_for_Intrinsic_Image_Decomposition_From_a_Single_Image_CVPR_2020_paper.html
AUTHORS: Yunfei Liu, Yu Li, Shaodi You, Feng Lu
HIGHLIGHT: In this paper, we propose a novel unsupervised intrinsic image decomposition framework, which relies on neither labeled training data nor hand-crafted priors.

322, TITLE: Multi-Domain Learning for Accurate and Few-Shot Color Constancy
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiao_Multi-Domain_Learning_for_Accurate_and_Few-Shot_Color_Constancy_CVPR_2020_paper.html
AUTHORS: Jin Xiao, Shuhang Gu, Lei Zhang
HIGHLIGHT: In this paper, we start a pioneer work to introduce multi-domain learning to color constancy area.

323, TITLE: PANDA: A Gigapixel-Level Human-Centric Video Dataset
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_PANDA_A_Gigapixel-Level_Human-Centric_Video_Dataset_CVPR_2020_paper.html
AUTHORS: Xueyang Wang, Xiya Zhang, Yinheng Zhu, Yuchen Guo, Xiaoyun Yuan, Liuyu Xiang, Zerun Wang, Guiguang Ding, David Brady, Qionghai Dai, Lu Fang
HIGHLIGHT: We present PANDA, the first gigaPixel-level humAN-centric viDeo dAtaset, for large-scale, long-term, and multi-object visual analysis.

324, TITLE: Cross-View Tracking for Multi-Human 3D Pose Estimation at Over 100 FPS
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Cross-View_Tracking_for_Multi-Human_3D_Pose_Estimation_at_Over_100_CVPR_2020_paper.html
AUTHORS: Long Chen, Haizhou Ai, Rui Chen, Zijie Zhuang, Shuang Liu
HIGHLIGHT: In this paper, we present a novel solution for multi-human 3D pose estimation from multiple calibrated camera views.

325, TITLE: Spatial-Temporal Graph Convolutional Network for Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Spatial-Temporal_Graph_Convolutional_Network_for_Video-Based_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Jinrui Yang, Wei-Shi Zheng, Qize Yang, Ying-Cong Chen, Qi Tian
HIGHLIGHT: In this work, we propose a novel Spatial-Temporal Graph Convolutional Network (STGCN) to solve these problems.

326, TITLE: Salience-Guided Cascaded Suppression Network for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Salience-Guided_Cascaded_Suppression_Network_for_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Xuesong Chen, Canmiao Fu, Yong Zhao, Feng Zheng, Jingkuan Song, Rongrong Ji, Yi Yang
HIGHLIGHT: To handle this limitation, we propose a novel Salience-guided Cascaded Suppression Network (SCSN) which enables the model to mine diverse salient features and integrate these features into the final representation by a cascaded manner.

327, TITLE: Fashion Outfit Complementary Item Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Fashion_Outfit_Complementary_Item_Retrieval_CVPR_2020_paper.html
AUTHORS: Yen-Liang Lin, Son Tran, Larry S. Davis
HIGHLIGHT: We propose a new framework for outfit complementary item retrieval.

328, TITLE: Learning Event-Based Motion Deblurring
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Learning_Event-Based_Motion_Deblurring_CVPR_2020_paper.html
AUTHORS: Zhe Jiang, Yu Zhang, Dongqing Zou, Jimmy Ren, Jiancheng Lv, Yebin Liu

HIGHLIGHT: In this paper, we start from a sequential formulation of event-based motion deblurring, then show how its optimization can be unfolded with a novel end-toend deep architecture.

329, TITLE: Domain Decluttering: Simplifying Images to Mitigate Synthetic-Real Domain Shift and Improve Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Domain_Decluttering_Simplifying_Images_to_Mitigate_Synthetic-Real_Domain_Shift_and_CVPR_2020_paper.html
AUTHORS: Yunhan Zhao, Shu Kong, Daeyun Shin, Charless Fowlkes
HIGHLIGHT: Based on these observations, we develop an attention module that learns to identify and remove difficult out-of-domain regions in real images in order to improve depth prediction for a model trained primarily on synthetic data.

330, TITLE: Neural Blind Deconvolution Using Deep Priors
http://openaccess.thecvf.com/content_CVPR_2020/html/Ren_Neural_Blind_Deconvolution_Using_Deep_Priors_CVPR_2020_paper.html
AUTHORS: Dongwei Ren, Kai Zhang, Qilong Wang, Qinghua Hu, Wangmeng Zuo
HIGHLIGHT: To connect MAP and deep models, we in this paper present two generative networks for respectively modeling the deep priors of clean image and blur kernel, and propose an unconstrained neural optimization solution to blind deconvolution.

331, TITLE: Anisotropic Convolutional Networks for 3D Semantic Scene Completion
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Anisotropic_Convolutional_Networks_for_3D_Semantic_Scene_Completion_CVPR_2020_paper.html
AUTHORS: Jie Li, Kai Han, Peng Wang, Yu Liu, Xia Yuan
HIGHLIGHT: To handle such variations, we propose a novel module called anisotropic convolution, which properties with flexibility and power impossible for the competing methods such as standard 3D convolution and some of its variations.

332, TITLE: TDAN: Temporally-Deformable Alignment Network for Video Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Tian_TDAN_Temporally-Deformable_Alignment_Network_for_Video_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Yapeng Tian, Yulun Zhang, Yun Fu, Chenliang Xu
HIGHLIGHT: To overcome the limitation, we propose a temporally-deformable alignment network (TDAN) to adaptively align the reference frame and each supporting frame at the feature level without computing optical flow.

333, TITLE: Zooming Slow-Mo: Fast and Accurate One-Stage Space-Time Video Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_Zooming_Slow-Mo_Fast_and_Accurate_One-Stage_Space-Time_Video_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Xiaoyu Xiang, Yapeng Tian, Yulun Zhang, Yun Fu, Jan P. Allebach, Chenliang Xu
HIGHLIGHT: In this paper, we explore the space-time video super-resolution task, which aims to generate a high-resolution (HR) slow-motion video from a low frame rate (LFR), low-resolution (LR) video.

334, TITLE: Fast MSER
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Fast_MSER_CVPR_2020_paper.html
AUTHORS: Hailiang Xu, Siqi Xie, Fan Chen
HIGHLIGHT: In this paper, we propose two novel MSER algorithms, called Fast MSER V1 and V2.

335, TITLE: Unsupervised Person Re-Identification via Softened Similarity Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Unsupervised_Person_Re-Identification_via_Softened_Similarity_Learning_CVPR_2020_paper.html
AUTHORS: Yutian Lin, Lingxi Xie, Yu Wu, Chenggang Yan, Qi Tian
HIGHLIGHT: In this paper, we follow the iterative training mechanism but discard clustering, since it incurs loss from hard quantization, yet its only product, image-level similarity, can be easily replaced by pairwise computation and a softened classification task.

336, TITLE: COCAS: A Large-Scale Clothes Changing Person Dataset for Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_COCAS_A_Large-Scale_Clothes_Changing_Person_Dataset_for_Re-Identification_CVPR_2020_paper.html
AUTHORS: Shijie Yu, Shihua Li, Dapeng Chen, Rui Zhao, Junjie Yan, Yu Qiao
HIGHLIGHT: To address the clothes changing person re-id problem, we construct a novel large-scale re-id benchmark named Clothes Changing Person Set (COCAS), which provides multiple images of the same identity with different clothes.

337, TITLE: Learning Formation of Physically-Based Face Attributes

http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Learning_Formation_of_Physically-
Based_Face_Attributes_CVPR_2020_paper.html
AUTHORS:　　　　Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha
Prasad, Bipin Kishore, Jun Xing, Hao Li
HIGHLIGHT:　　　　Based on a combined data set of 4000 high resolution facial scans, we introduce a non-linear morphable face
model, capable of producing multifarious face geometry of pore-level resolution, coupled with material attributes for use in
physically-based rendering.

338, TITLE:　　　　Generalized Product Quantization Network for Semi-Supervised Image Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Jang_Generalized_Product_Quantization_Network_for_Semi-
Supervised_Image_Retrieval_CVPR_2020_paper.html
AUTHORS:　　　　Young Kyun Jang, Nam Ik Cho
HIGHLIGHT:　　　　To resolve this issue, we propose the first quantization-based semi-supervised image retrieval scheme:
Generalized Product Quantization (GPQ) network.

339, TITLE:　　　　Stereoscopic Flash and No-Flash Photography for Shape and Albedo Recovery
http://openaccess.thecvf.com/content_CVPR_2020/html/Cao_Stereoscopic_Flash_and_No-
Flash_Photography_for_Shape_and_Albedo_Recovery_CVPR_2020_paper.html
AUTHORS:　　　　Xu Cao, Michael Waechter, Boxin Shi, Ye Gao, Bo Zheng, Yasuyuki Matsushita
HIGHLIGHT:　　　　We present a minimal imaging setup that harnesses both geometric and photometric approaches for shape and
albedo recovery.

340, TITLE:　　　　Context-Aware Group Captioning via Self-Attention and Contrastive Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Context-Aware_Group_Captioning_via_Self-
Attention_and_Contrastive_Features_CVPR_2020_paper.html
AUTHORS:　　　　Zhuowan Li, Quan Tran, Long Mai, Zhe Lin, Alan L. Yuille
HIGHLIGHT:　　　　To solve this problem, we propose a framework combining self-attention mechanism with contrastive feature
construction to effectively summarize common information from each image group while capturing discriminative information
between them.

341, TITLE:　　　　MEBOW: Monocular Estimation of Body Orientation in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_MEBOW_Monocular_Estimation_of_Body_Orientation_in_the_Wild_
CVPR_2020_paper.html
AUTHORS:　　　　Chenyan Wu, Yukun Chen, Jiajia Luo, Che-Chun Su, Anuja Dawane, Bikramjot Hanzra, Zhuo Deng, Bilan Liu,
James Z. Wang, Cheng-hao Kuo
HIGHLIGHT:　　　　We present COCO-MEBOW (Monocular Estimation of Body Orientation in the Wild), a new large-scale
dataset for orientation estimation from a single in-the-wild image.

342, TITLE:　　　　Distilling Image Dehazing With Heterogeneous Task Imitation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hong_Distilling_Image_Dehazing_With_Heterogeneous_Task_Imitation_C
VPR_2020_paper.html
AUTHORS:　　　　Ming Hong, Yuan Xie, Cuihua Li, Yanyun Qu
HIGHLIGHT:　　　　In this paper, we propose a knowledge-distill dehazing network which distills image dehazing with the
heterogeneous task imitation.

343, TITLE:　　　　Select, Supplement and Focus for RGB-D Saliency Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Select_Supplement_and_Focus_for_RGB-
D_Saliency_Detection_CVPR_2020_paper.html
AUTHORS:　　　　Miao Zhang, Weisong Ren, Yongri Piao, Zhengkun Rong, Huchuan Lu
HIGHLIGHT:　　　　In this paper, we propose a new framework for accurate RGB-D saliency detection taking account of local and
global complementarities from two modalities.

344, TITLE:　　　　Transfer Learning From Synthetic to Real-Noise Denoising With Adaptive Instance Normalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Transfer_Learning_From_Synthetic_to_Real-
Noise_Denoising_With_Adaptive_Instance_CVPR_2020_paper.html
AUTHORS:　　　　Yoonsik Kim, Jae Woong Soh, Gu Yong Park, Nam Ik Cho
HIGHLIGHT:　　　　In order to cope with various and complex real-noise, we propose a well-generalized denoising architecture and
a transfer learning scheme.

345, TITLE:　　　　On Joint Estimation of Pose, Geometry and svBRDF From a Handheld Scanner
http://openaccess.thecvf.com/content_CVPR_2020/html/Schmitt_On_Joint_Estimation_of_Pose_Geometry_and_svBRDF_From_a_C
VPR_2020_paper.html

AUTHORS: Carolin Schmitt, Simon Donne, Gernot Riegler, Vladlen Koltun, Andreas Geiger
HIGHLIGHT: We propose a novel formulation for joint recovery of camera pose, object geometry and spatially-varying BRDF.

346, TITLE: Differentiable Volumetric Rendering: Learning Implicit 3D Representations Without 3D Supervision
http://openaccess.thecvf.com/content_CVPR_2020/html/Niemeyer_Differentiable_Volumetric_Rendering_Learning_Implicit_3D_Representations_Without_3D_Supervision_CVPR_2020_paper.html
AUTHORS: Michael Niemeyer, Lars Mescheder, Michael Oechsle, Andreas Geiger
HIGHLIGHT: In this work, we propose a differentiable rendering formulation for implicit shape and texture representations.

347, TITLE: Meta-Transfer Learning for Zero-Shot Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Soh_Meta-Transfer_Learning_for_Zero-Shot_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Jae Woong Soh, Sunwoo Cho, Nam Ik Cho
HIGHLIGHT: In this paper, we present Meta-Transfer Learning for Zero-Shot Super-Resolution (MZSR), which leverages ZSSR.

348, TITLE: Solving Jigsaw Puzzles With Eroded Boundaries
http://openaccess.thecvf.com/content_CVPR_2020/html/Bridger_Solving_Jigsaw_Puzzles_With_Eroded_Boundaries_CVPR_2020_paper.html
AUTHORS: Dov Bridger, Dov Danon, Ayellet Tal
HIGHLIGHT: This paper focuses on a specific variant of the problem--solving puzzles with eroded boundaries.

349, TITLE: Context-Aware Attention Network for Image-Text Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Context-Aware_Attention_Network_for_Image-Text_Retrieval_CVPR_2020_paper.html
AUTHORS: Qi Zhang, Zhen Lei, Zhaoxiang Zhang, Stan Z. Li
HIGHLIGHT: In this work, we propose a unified Context-Aware Attention Network (CAAN), which selectively focuses on critical local fragments (regions and words) by aggregating the global context.

350, TITLE: M-LVC: Multiple Frames Prediction for Learned Video Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_M-LVC_Multiple_Frames_Prediction_for_Learned_Video_Compression_CVPR_2020_paper.html
AUTHORS: Jianping Lin, Dong Liu, Houqiang Li, Feng Wu
HIGHLIGHT: We propose an end-to-end learned video compression scheme for low-latency scenarios.

351, TITLE: Efficient Dynamic Scene Deblurring Using Spatially Variant Deconvolution Network With Optical Flow Guided Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Yuan_Efficient_Dynamic_Scene_Deblurring_Using_Spatially_Variant_Deconvolution_Network_With_CVPR_2020_paper.html
AUTHORS: Yuan Yuan, Wei Su, Dandan Ma
HIGHLIGHT: In this paper, we start from the deblurring deconvolution operation, then design an effective and real-time deblurring network.

352, TITLE: Single Image Reflection Removal Through Cascaded Refinement
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Single_Image_Reflection_Removal_Through_Cascaded_Refinement_CVPR_2020_paper.html
AUTHORS: Chao Li, Yixiao Yang, Kun He, Stephen Lin, John E. Hopcroft
HIGHLIGHT: Inspired by iterative structure reduction for hidden community detection in social networks, we propose an Iterative Boost Convolutional LSTM Network (IBCLN) that enables cascaded prediction for reflection removal.

353, TITLE: From Patches to Pictures (PaQ-2-PiQ): Mapping the Perceptual Space of Picture Quality
http://openaccess.thecvf.com/content_CVPR_2020/html/Ying_From_Patches_to_Pictures_PaQ-2-PiQ_Mapping_the_Perceptual_Space_of_CVPR_2020_paper.html
AUTHORS: Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, Alan Bovik
HIGHLIGHT: To advance progress on this problem, we introduce the largest (by far) subjective picture quality database, containing about 40, 000 real-world distorted pictures and 120, 000 patches, on which we collected about 4M human judgments of picture quality.

354, TITLE: Video to Events: Recycling Video Datasets for Event Cameras

http://openaccess.thecvf.com/content_CVPR_2020/html/Gehrig_Video_to_Events_Recycling_Video_Datasets_for_Event_Cameras_CVPR_2020_paper.html
AUTHORS:	Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrio, Davide Scaramuzza
HIGHLIGHT:	In this paper, we present a method that addresses these needs by converting any existing video dataset recorded with conventional cameras to synthetic event data.

355, TITLE:	Composed Query Image Retrieval Using Locally Bounded Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Hosseinzadeh_Composed_Query_Image_Retrieval_Using_Locally_Bounded_Features_CVPR_2020_paper.html
AUTHORS:	Mehrdad Hosseinzadeh, Yang Wang
HIGHLIGHT:	In this paper, we propose a novel method that represents the image using a set of local areas in the image.

356, TITLE:	Spatially-Attentive Patch-Hierarchical Network for Adaptive Motion Deblurring
http://openaccess.thecvf.com/content_CVPR_2020/html/Suin_Spatially-Attentive_Patch-Hierarchical_Network_for_Adaptive_Motion_Deblurring_CVPR_2020_paper.html
AUTHORS:	Maitreya Suin, Kuldeep Purohit, A. N. Rajagopalan
HIGHLIGHT:	In this work, we propose an efficient pixel adaptive and feature attentive design for handling large blur variations across different spatial locations and process each test image adaptively.

357, TITLE:	End-to-End Illuminant Estimation Based on Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_End-to-End_Illuminant_Estimation_Based_on_Deep_Metric_Learning_CVPR_2020_paper.html
AUTHORS:	Bolei Xu, Jingxin Liu, Xianxu Hou, Bozhi Liu, Guoping Qiu
HIGHLIGHT:	To overcome this problem, we introduce a deep metric learning approach named Illuminant-Guided Triplet Network (IGTN) to color constancy.

358, TITLE:	Variational-EM-Based Deep Learning for Noise-Blind Image Deblurring
http://openaccess.thecvf.com/content_CVPR_2020/html/Nan_Variational-EM-Based_Deep_Learning_for_Noise-Blind_Image_Deblurring_CVPR_2020_paper.html
AUTHORS:	Yuesong Nan, Yuhui Quan, Hui Ji
HIGHLIGHT:	This paper aims at developing a deep learning framework for deblurring images with unknown noise level.

359, TITLE:	Image Demoireing with Learnable Bandpass Filters
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Image_Demoireing_with_Learnable_Bandpass_Filters_CVPR_2020_paper.html
AUTHORS:	Bolun Zheng, Shanxin Yuan, Gregory Slabaugh, Ales Leonardis
HIGHLIGHT:	In this paper, we propose a novel multiscale bandpass convolutional neural network (MBCNN) to address this problem.

360, TITLE:	Assessing Image Quality Issues for Real-World Problems
http://openaccess.thecvf.com/content_CVPR_2020/html/Chiu_Assessing_Image_Quality_Issues_for_Real-World_Problems_CVPR_2020_paper.html
AUTHORS:	Tai-Yin Chiu, Yinan Zhao, Danna Gurari
HIGHLIGHT:	We introduce a new large-scale dataset that links the assessment of image quality issues to two practical vision tasks: image captioning and visual question answering.

361, TITLE:	Memory-Efficient Hierarchical Neural Architecture Search for Image Denoising
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Memory-Efficient_Hierarchical_Neural_Architecture_Search_for_Image_Denoising_CVPR_2020_paper.html
AUTHORS:	Haokui Zhang, Ying Li, Hao Chen, Chunhua Shen
HIGHLIGHT:	In this paper, we propose HiNAS (Hierarchical NAS), an effort towards employing NAS to automatically design effective neural network architectures for image denoising.

362, TITLE:	Blindly Assess Image Quality in the Wild Guided by a Self-Adaptive Hyper Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Su_Blindly_Assess_Image_Quality_in_the_Wild_Guided_by_a_CVPR_2020_paper.html
AUTHORS:	Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, Yanning Zhang
HIGHLIGHT:	To deal with the challenge, we propose a self-adaptive hyper network architecture to blind assess image quality in the wild.

363, TITLE:	Perceptual Quality Assessment of Smartphone Photography

http://openaccess.thecvf.com/content_CVPR_2020/html/Fang_Perceptual_Quality_Assessment_of_Smartphone_Photography_CVPR_2020_paper.html
AUTHORS: Yuming Fang, Hanwei Zhu, Yan Zeng, Kede Ma, Zhou Wang
HIGHLIGHT: We introduce the Smartphone Photography Attribute and Quality (SPAQ) database, consisting of 11,125 pictures taken by 66 smartphones, where each image is attached with so far the richest annotations.

364, TITLE: Don't Hit Me! Glass Detection in Real-World Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Mei_Dont_Hit_Me_Glass_Detection_in_Real-World_Scenes_CVPR_2020_paper.html
AUTHORS: Haiyang Mei, Xin Yang, Yang Wang, Yuanyuan Liu, Shengfeng He, Qiang Zhang, Xiaopeng Wei, Rynson W.H. Lau
HIGHLIGHT: In this paper, we propose an important problem of detecting glass from a single RGB image.

365, TITLE: Progressive Mirror Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Progressive_Mirror_Detection_CVPR_2020_paper.html
AUTHORS: Jiaying Lin, Guodong Wang, Rynson W.H. Lau
HIGHLIGHT: Hence, we propose a model in this paper to progressively learn the content similarity between the inside and outside of the mirror while explicitly detecting the mirror edges.

366, TITLE: Category-Level Articulated Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Category-Level_Articulated_Object_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Xiaolong Li, He Wang, Li Yi, Leonidas J. Guibas, A. Lynn Abbott, Shuran Song
HIGHLIGHT: We present a novel category-level approach that correctly accommodates object instances previously unseen during training.

367, TITLE: Unbiased Scene Graph Generation From Biased Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Unbiased_Scene_Graph_Generation_From_Biased_Training_CVPR_2020_paper.html
AUTHORS: Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, Hanwang Zhang
HIGHLIGHT: In this paper, we present a novel SGG framework based on causal inference but not the conventional likelihood.

368, TITLE: Dynamic Graph Message Passing Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Dynamic_Graph_Message_Passing_Networks_CVPR_2020_paper.html
AUTHORS: Li Zhang, Dan Xu, Anurag Arnab, Philip H.S. Torr
HIGHLIGHT: We propose a dynamic graph message passing network, that significantly reduces the computational complexity compared to related works modelling a fully-connected graph.

369, TITLE: Weakly Supervised Visual Semantic Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Zareian_Weakly_Supervised_Visual_Semantic_Parsing_CVPR_2020_paper.html
AUTHORS: Alireza Zareian, Svebor Karaman, Shih-Fu Chang
HIGHLIGHT: In this paper, we address those two limitations by first proposing a generalized formulation of SGG, namely Visual Semantic Parsing, which disentangles entity and predicate recognition, and enables sub-quadratic performance.

370, TITLE: GPS-Net: Graph Property Sensing Network for Scene Graph Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_GPS-Net_Graph_Property_Sensing_Network_for_Scene_Graph_Generation_CVPR_2020_paper.html
AUTHORS: Xin Lin, Changxing Ding, Jinquan Zeng, Dacheng Tao
HIGHLIGHT: Accordingly, in this paper, we propose a Graph Property Sensing Network (GPS-Net) that fully explores these three properties for SGG.

371, TITLE: End-to-End Optimization of Scene Layout
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_End-to-End_Optimization_of_Scene_Layout_CVPR_2020_paper.html
AUTHORS: Andrew Luo, Zhoutong Zhang, Jiajun Wu, Joshua B. Tenenbaum
HIGHLIGHT: We propose an end-to-end variational generative model for scene layout synthesis conditioned on scene graphs.

372, TITLE: Unsupervised Intra-Domain Adaptation for Semantic Segmentation Through Self-Supervision
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Unsupervised_Intra-Domain_Adaptation_for_Semantic_Segmentation_Through_Self-Supervision_CVPR_2020_paper.html

AUTHORS:　　　　Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, In So Kweon
HIGHLIGHT:　　　　In this work, we propose a two-step self-supervised domain adaptation approach to minimize the inter-domain and intra-domain gap together.

373, TITLE:　　　　Dual Super-Resolution Learning for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Dual_Super-Resolution_Learning_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS:　　　　Li Wang, Dong Li, Yousong Zhu, Lu Tian, Yi Shan
HIGHLIGHT:　　　　In this paper, we propose a simple and flexible two-stream framework named Dual Super-Resolution Learning (DSRL) to effectively improve the segmentation accuracy without introducing extra computation costs.

374, TITLE:　　　　Self-Supervised Scene De-Occlusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhan_Self-Supervised_Scene_De-Occlusion_CVPR_2020_paper.html
AUTHORS:　　　　Xiaohang Zhan, Xingang Pan, Bo Dai, Ziwei Liu, Dahua Lin, Chen Change Loy
HIGHLIGHT:　　　　In this paper, we investigate the problem of scene de-occlusion, which aims to recover the underlying occlusion ordering and complete the invisible parts of occluded objects.

375, TITLE:　　　　BANet: Bidirectional Aggregation Network With Occlusion Handling for Panoptic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_BANet_Bidirectional_Aggregation_Network_With_Occlusion_Handling_for_Panoptic_Segmentation_CVPR_2020_paper.html
AUTHORS:　　　　Yifeng Chen, Guangchen Lin, Songyuan Li, Omar Bourahla, Yiming Wu, Fangfang Wang, Junyi Feng, Mingliang Xu, Xi Li
HIGHLIGHT:　　　　Motivated by these observations, we propose a novel deep panoptic segmentation scheme based on a bidirectional learning pipeline.

376, TITLE:　　　　CPR-GCN: Conditional Partial-Residual Graph Convolutional Network in Automated Anatomical Labeling of Coronary Arteries
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_CPR-GCN_Conditional_Partial-Residual_Graph_Convolutional_Network_in_Automated_Anatomical_Labeling_CVPR_2020_paper.html
AUTHORS:　　　　Han Yang, Xingjian Zhen, Ying Chi, Lei Zhang, Xian-Sheng Hua
HIGHLIGHT:　　　　Motivated by the wide application of the graph neural network in structured data, in this paper, we propose a conditional partial-residual graph convolutional network (CPR-GCN), which takes both position and CT image into consideration, since CT image contains abundant information such as branch size and spanning direction.

377, TITLE:　　　　Cross-View Correspondence Reasoning Based on Bipartite Graph Convolutional Network for Mammogram Mass Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Cross-View_Correspondence_Reasoning_Based_on_Bipartite_Graph_Convolutional_Network_for_CVPR_2020_paper.html
AUTHORS:　　　　Yuhang Liu, Fandong Zhang, Qianyi Zhang, Siwen Wang, Yizhou Wang, Yizhou Yu
HIGHLIGHT:　　　　In this paper, we introduce bipartite graph convolutional network to endow existing methods with cross-view reasoning ability of radiologists in mammogram mass detection.

378, TITLE:　　　　MPM: Joint Representation of Motion and Position Map for Cell Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Hayashida_MPM_Joint_Representation_of_Motion_and_Position_Map_for_Cell_CVPR_2020_paper.html
AUTHORS:　　　　Junya Hayashida, Kazuya Nishimura, Ryoma Bise
HIGHLIGHT:　　　　In this paper, we propose the Motion and Position Map (MPM) that jointly represents both detection and association for not only migration but also cell division.

379, TITLE:　　　　Deep Distance Transform for Tubular Structure Segmentation in CT Scans
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Deep_Distance_Transform_for_Tubular_Structure_Segmentation_in_CT_Scans_CVPR_2020_paper.html
AUTHORS:　　　　Yan Wang, Xu Wei, Fengze Liu, Jieneng Chen, Yuyin Zhou, Wei Shen, Elliot K. Fishman, Alan L. Yuille
HIGHLIGHT:　　　　Inspired by this, we propose a geometry-aware tubular structure segmentation method, Deep Distance Transform (DDT), which combines intuitions from the classical distance transform for skeletonization and modern deep segmentation networks.

380, TITLE:　　　　Instance Segmentation of Biological Images Using Harmonic Embeddings
http://openaccess.thecvf.com/content_CVPR_2020/html/Kulikov_Instance_Segmentation_of_Biological_Images_Using_Harmonic_Embeddings_CVPR_2020_paper.html
AUTHORS:　　　　Victor Kulikov, Victor Lempitsky

HIGHLIGHT: We present a new instance segmentation approach tailored to biological images, where instances may correspond to individual cells, organisms or plant parts.

381, TITLE: Multi-scale Domain-adversarial Multiple-instance CNN for Cancer Subtype Classification with Unannotated Histopathological Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Hashimoto_Multi-scale_Domain-adversarial_Multiple-instance_CNN_for_Cancer_Subtype_Classification_with_Unannotated_CVPR_2020_paper.html
AUTHORS: Noriaki Hashimoto, Daisuke Fukushima, Ryoichi Koga, Yusuke Takagi, Kaho Ko, Kei Kohno, Masato Nakaguro, Shigeo Nakamura, Hidekata Hontani, Ichiro Takeuchi
HIGHLIGHT: We propose a new method for cancer subtype classification from histopathological images, which can automatically detect tumor-specific features in a given whole slide image (WSI).

382, TITLE: SOS: Selective Objective Switch for Rapid Immunofluorescence Whole Slide Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Maksoud_SOS_Selective_Objective_Switch_for_Rapid_Immunofluorescence_Whole_Slide_Image_CVPR_2020_paper.html
AUTHORS: Sam Maksoud, Kun Zhao, Peter Hobson, Anthony Jennings, Brian C. Lovell
HIGHLIGHT: In this paper, we demonstrate that conventional patch-based processing is redundant for certain WSI classification tasks where high resolution is only required in a minority of cases.

383, TITLE: Task Agnostic Robust Learning on Corrupt Outputs by Correlation-Guided Mixture Density Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Choi_Task_Agnostic_Robust_Learning_on_Corrupt_Outputs_by_Correlation-Guided_Mixture_CVPR_2020_paper.html
AUTHORS: Sungjoon Choi, Sanghoon Hong, Kyungjae Lee, Sungbin Lim
HIGHLIGHT: In this paper, we focus on weakly supervised learning with noisy training data for both classification and regression problems.

384, TITLE: METAL: Minimum Effort Temporal Activity Localization in Untrimmed Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_METAL_Minimum_Effort_Temporal_Activity_Localization_in_Untrimmed_Videos_CVPR_2020_paper.html
AUTHORS: Da Zhang, Xiyang Dai, Yuan-Fang Wang
HIGHLIGHT: Towards this objective, we propose a novel Similarity Pyramid Network (SPN) that adopts the few-shot learning technique of Relation Network and directly encodes hierarchical multi-scale correlations, which we learn by optimizing two complimentary loss functions in an end-to-end manner.

385, TITLE: Neural Data Server: A Large-Scale Search Engine for Transfer Learning Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Neural_Data_Server_A_Large-Scale_Search_Engine_for_Transfer_Learning_CVPR_2020_paper.html
AUTHORS: Xi Yan, David Acuna, Sanja Fidler
HIGHLIGHT: We introduce Neural Data Server (NDS), a large-scale search engine for finding the most useful transfer learning data to the target domain.

386, TITLE: Revisiting Knowledge Distillation via Label Smoothing Regularization
http://openaccess.thecvf.com/content_CVPR_2020/html/Yuan_Revisiting_Knowledge_Distillation_via_Label_Smoothing_Regularization_CVPR_2020_paper.html
AUTHORS: Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, Jiashi Feng
HIGHLIGHT: In this work, we challenge this common belief by following experimental observations: 1) beyond the acknowledgment that the teacher can improve the student, the student can also enhance the teacher significantly by reversing the KD procedure; 2) a poorly-trained teacher with much lower accuracy than the student can still improve the latter significantly.

387, TITLE: WCP: Worst-Case Perturbations for Semi-Supervised Deep Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_WCP_Worst-Case_Perturbations_for_Semi-Supervised_Deep_Learning_CVPR_2020_paper.html
AUTHORS: Liheng Zhang, Guo-Jun Qi
HIGHLIGHT: In this paper, we present a novel regularization mechanism for training deep networks by minimizing the Worse-Case Perturbation (WCP).

388, TITLE: DEPARA: Deep Attribution Graph for Deep Knowledge Transferability
http://openaccess.thecvf.com/content_CVPR_2020/html/Song_DEPARA_Deep_Attribution_Graph_for_Deep_Knowledge_Transferability_CVPR_2020_paper.html
AUTHORS: Jie Song, Yixin Chen, Jingwen Ye, Xinchao Wang, Chengchao Shen, Feng Mao, Mingli Song
HIGHLIGHT: In this paper, we propose the DEeP Attribution gRAph (DEPARA) to investigate the transferability of knowledge learned from PR-DNNs.

389, TITLE: Conditional Channel Gated Networks for Task-Aware Continual Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Abati_Conditional_Channel_Gated_Networks_for_Task-Aware_Continual_Learning_CVPR_2020_paper.html
AUTHORS: Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, Babak Ehteshami Bejnordi
HIGHLIGHT: In this work, we introduce a novel framework to tackle this problem with conditional computation.

390, TITLE: Towards Discriminability and Diversity: Batch Nuclear-Norm Maximization Under Label Insufficient Situations
http://openaccess.thecvf.com/content_CVPR_2020/html/Cui_Towards_Discriminability_and_Diversity_Batch_Nuclear-Norm_Maximization_Under_Label_Insufficient_CVPR_2020_paper.html
AUTHORS: Shuhao Cui, Shuhui Wang, Junbao Zhuo, Liang Li, Qingming Huang, Qi Tian
HIGHLIGHT: Accordingly, to improve both discriminability and diversity, we propose Batch Nuclear-norm Maximization (BNM) on the output matrix.

391, TITLE: FocalMix: Semi-Supervised Learning for 3D Medical Image Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_FocalMix_Semi-Supervised_Learning_for_3D_Medical_Image_Detection_CVPR_2020_paper.html
AUTHORS: Dong Wang, Yuan Zhang, Kexin Zhang, Liwei Wang
HIGHLIGHT: In this paper, we propose a novel method, called FocalMix, which, to the best of our knowledge, is the first to leverage recent advances in semi-supervised learning (SSL) for 3D medical image detection.

392, TITLE: Learning 3D Semantic Scene Graphs From 3D Indoor Reconstructions
http://openaccess.thecvf.com/content_CVPR_2020/html/Wald_Learning_3D_Semantic_Scene_Graphs_From_3D_Indoor_Reconstructions_CVPR_2020_paper.html
AUTHORS: Johanna Wald, Helisa Dhamo, Nassir Navab, Federico Tombari
HIGHLIGHT: In our work we focus on scene graphs, a data structure that organizes the entities of a scene in a graph, where objects are nodes and their relationships modeled as edges.

393, TITLE: Self-Supervised Viewpoint Learning From Image Collections
http://openaccess.thecvf.com/content_CVPR_2020/html/Mustikovela_Self-Supervised_Viewpoint_Learning_From_Image_Collections_CVPR_2020_paper.html
AUTHORS: Siva Karthik Mustikovela, Varun Jampani, Shalini De Mello, Sifei Liu, Umar Iqbal, Carsten Rother, Jan Kautz
HIGHLIGHT: We propose a novel learning framework which incorporates an analysis-by-synthesis paradigm to reconstruct images in a viewpoint aware manner with a generative network, along with symmetry and adversarial constraints to successfully supervise our viewpoint estimation network.

394, TITLE: Two-Shot Spatially-Varying BRDF and Shape Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Boss_Two-Shot_Spatially-Varying_BRDF_and_Shape_Estimation_CVPR_2020_paper.html
AUTHORS: Mark Boss, Varun Jampani, Kihwan Kim, Hendrik P.A. Lensch, Jan Kautz
HIGHLIGHT: We propose a novel deep learning architecture with a stage-wise estimation of shape and SVBRDF.

395, TITLE: Variational Context-Deformable ConvNets for Indoor Scene Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiong_Variational_Context-Deformable_ConvNets_for_Indoor_Scene_Parsing_CVPR_2020_paper.html
AUTHORS: Zhitong Xiong, Yuan Yuan, Nianhui Guo, Qi Wang
HIGHLIGHT: Thus, in this paper, we propose a novel variational context-deformable (VCD) module to learn adaptive receptive-field in a structured fashion.

396, TITLE: Strip Pooling: Rethinking Spatial Pooling for Scene Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Hou_Strip_Pooling_Rethinking_Spatial_Pooling_for_Scene_Parsing_CVPR_2020_paper.html
AUTHORS: Qibin Hou, Li Zhang, Ming-Ming Cheng, Jiashi Feng
HIGHLIGHT: In this paper, beyond conventional spatial pooling that usually has a regular shape of NxN, we rethink the formulation of spatial pooling by introducing a new pooling strategy, called strip pooling, which considers a long but narrow kernel, i.e., 1xN or Nx1.

397, TITLE: Few-Shot Object Detection With Attention-RPN and Multi-Relation Detector

http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_Few-Shot_Object_Detection_With_Attention-RPN_and_Multi-Relation_Detector_CVPR_2020_paper.html
AUTHORS: Qi Fan, Wei Zhuo, Chi-Keung Tang, Yu-Wing Tai
HIGHLIGHT: In this paper, we propose a novel few-shot object detection network that aims at detecting objects of unseen categories with only a few annotated examples.

398, TITLE: What Can Be Transferred: Unsupervised Domain Adaptation for Endoscopic Lesions Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_What_Can_Be_Transferred_Unsupervised_Domain_Adaptation_for_Endoscopic_Lesions_CVPR_2020_paper.html
AUTHORS: Jiahua Dong, Yang Cong, Gan Sun, Bineng Zhong, Xiaowei Xu
HIGHLIGHT: To address these challenges, we develop a new unsupervised semantic transfer model including two complementary modules (i.e., T_D and T_F) for endoscopic lesions segmentation, which can alternatively determine where and how to explore transferable domain-invariant knowledge between labeled source lesions dataset (e.g., gastroscope) and unlabeled target diseases dataset (e.g., enteroscopy).

399, TITLE: ADINet: Attribute Driven Incremental Network for Retinal Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Meng_ADINet_Attribute_Driven_Incremental_Network_for_Retinal_Image_Classification_CVPR_2020_paper.html
AUTHORS: Qier Meng, Satoh Shin'ichi
HIGHLIGHT: In this paper, we design a framework named "Attribute Driven Incremental Network" (ADINet), a new architecture that integrates class label prediction and attribute prediction into an incremental learning framework to boost the classification performance.

400, TITLE: Unsupervised Domain Adaptation With Hierarchical Gradient Synchronization
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Unsupervised_Domain_Adaptation_With_Hierarchical_Gradient_Synchronization_CVPR_2020_paper.html
AUTHORS: Lanqing Hu, Meina Kan, Shiguang Shan, Xilin Chen
HIGHLIGHT: Inspired by this, we propose a novel method called Hierarchical Gradient Synchronization to model the synchronization relationship among the local distribution pieces and global distribution, aiming for more precise domain-invariant features.

401, TITLE: Deep Grouping Model for Unified Perceptual Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Deep_Grouping_Model_for_Unified_Perceptual_Parsing_CVPR_2020_paper.html
AUTHORS: Zhiheng Li, Wenxuan Bao, Jiayang Zheng, Chenliang Xu
HIGHLIGHT: Overcoming these challenges, we propose a deep grouping model (DGM) that tightly marries the two types of representations and defines a bottom-up and a top-down process for feature exchanging.

402, TITLE: Where Am I Looking At? Joint Location and Orientation Estimation by Cross-View Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Where_Am_I_Looking_At_Joint_Location_and_Orientation_Estimation_CVPR_2020_paper.html
AUTHORS: Yujiao Shi, Xin Yu, Dylan Campbell, Hongdong Li
HIGHLIGHT: Therefore, we design a Dynamic Similarity Matching network to estimate cross-view orientation alignment during localization.

403, TITLE: Gum-Net: Unsupervised Geometric Matching for Fast and Accurate 3D Subtomogram Image Alignment and Averaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_Gum-Net_Unsupervised_Geometric_Matching_for_Fast_and_Accurate_3D_Subtomogram_CVPR_2020_paper.html
AUTHORS: Xiangrui Zeng, Min Xu
HIGHLIGHT: We propose a Geometric unsupervised matching Net-work (Gum-Net) for finding the geometric correspondence between two images with application to 3D subtomogram alignment and averaging.

404, TITLE: FDA: Fourier Domain Adaptation for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_FDA_Fourier_Domain_Adaptation_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Yanchao Yang, Stefano Soatto
HIGHLIGHT: We describe a simple method for unsupervised domain adaptation, whereby the discrepancy between the source and target distributions is reduced by swapping the low-frequency spectrum of one with the other.

405, TITLE: Foreground-Aware Relation Network for Geospatial Object Segmentation in High Spatial Resolution Remote Sensing Imagery

http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Foreground-Aware_Relation_Network_for_Geospatial_Object_Segmentation_in_High_Spatial_CVPR_2020_paper.html
AUTHORS: Zhuo Zheng, Yanfei Zhong, Junjue Wang, Ailong Ma
HIGHLIGHT: In this paper, we argue that the problems lie on the lack of foreground modeling and propose a foreground-aware relation network (FarSeg) from the perspectives of relation-based and optimization-based foreground modeling, to alleviate the above two problems.

406, TITLE: When2com: Multi-Agent Perception via Communication Graph Grouping
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_When2com_Multi-Agent_Perception_via_Communication_Graph_Grouping_CVPR_2020_paper.html
AUTHORS: Yen-Cheng Liu, Junjiao Tian, Nathaniel Glaser, Zsolt Kira
HIGHLIGHT: In this paper, we address the collaborative perception problem, where one agent is required to perform a perception task and can communicate and share information with other agents on the same task.

407, TITLE: Learning Human-Object Interaction Detection Using Interaction Points
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Learning_Human-Object_Interaction_Detection_Using_Interaction_Points_CVPR_2020_paper.html
AUTHORS: Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, Jian Sun
HIGHLIGHT: In this paper, we therefore propose a novel fully-convolutional approach that directly detects the interactions between human-object pairs.

408, TITLE: C2FNAS: Coarse-to-Fine Neural Architecture Search for 3D Medical Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_C2FNAS_Coarse-to-Fine_Neural_Architecture_Search_for_3D_Medical_Image_Segmentation_CVPR_2020_paper.html
AUTHORS: Qihang Yu, Dong Yang, Holger Roth, Yutong Bai, Yixiao Zhang, Alan L. Yuille, Daguang Xu
HIGHLIGHT: In this paper, we propose a coarse-to-fine neural architecture search (C2FNAS) to automatically search a 3D segmentation network from scratch without inconsistency on network size or input size.

409, TITLE: Adaptive Subspaces for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Simon_Adaptive_Subspaces_for_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Christian Simon, Piotr Koniusz, Richard Nock, Mehrtash Harandi
HIGHLIGHT: In this paper, we provide a framework for few-shot learning by introducing dynamic classifiers that are constructed from few samples.

410, TITLE: Learning to Detect Important People in Unlabelled Images for Semi-Supervised Important People Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Hong_Learning_to_Detect_Important_People_in_Unlabelled_Images_for_Semi-Supervised_CVPR_2020_paper.html
AUTHORS: Fa-Ting Hong, Wei-Hong Li, Wei-Shi Zheng
HIGHLIGHT: To overcome this problem, we propose learning important people detection on partially annotated images.

411, TITLE: Stochastic Sparse Subspace Clustering
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Stochastic_Sparse_Subspace_Clustering_CVPR_2020_paper.html
AUTHORS: Ying Chen, Chun-Guang Li, Chong You
HIGHLIGHT: In particular, we show that dropout is equivalent to adding a squared l_2 norm regularization on the representation coefficients, therefore induces denser solutions. Then, we reformulate the optimization problem as a consensus problem over a set of small-scale subproblems.

412, TITLE: CRNet: Cross-Reference Networks for Few-Shot Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_CRNet_Cross-Reference_Networks_for_Few-Shot_Segmentation_CVPR_2020_paper.html
AUTHORS: Weide Liu, Chi Zhang, Guosheng Lin, Fayao Liu
HIGHLIGHT: In this paper, we propose a cross-reference network (CRNet) for few-shot segmentation.

413, TITLE: Shoestring: Graph-Based Semi-Supervised Classification With Severely Limited Labeled Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Shoestring_Graph-Based_Semi-Supervised_Classification_With_Severely_Limited_Labeled_Data_CVPR_2020_paper.html
AUTHORS: Wanyu Lin, Zhaolin Gao, Baochun Li
HIGHLIGHT: To address the problem of semi-supervised learning in the presence of severely limited labeled samples, we propose a new framework, called Shoestring , that incorporates metric learning into the paradigm of graph-based semi-supervised learning.

414, TITLE: Uninformed Students: Student-Teacher Anomaly Detection With Discriminative Latent Embeddings
http://openaccess.thecvf.com/content_CVPR_2020/html/Bergmann_Uninformed_Students_Student-Teacher_Anomaly_Detection_With_Discriminative_Latent_Embeddings_CVPR_2020_paper.html
AUTHORS: Paul Bergmann, Michael Fauser, David Sattlegger, Carsten Steger
HIGHLIGHT: We introduce a powerful student-teacher framework for the challenging problem of unsupervised anomaly detection and pixel-precise anomaly segmentation in high-resolution images.

415, TITLE: 3D Sketch-Aware Semantic Scene Completion via Semi-Supervised Structure Prior
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_3D_Sketch-Aware_Semantic_Scene_Completion_via_Semi-Supervised_Structure_Prior_CVPR_2020_paper.html
AUTHORS: Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, Hongsheng Li
HIGHLIGHT: In this paper, we propose to devise a new geometry-based strategy to embed depth information with low-resolution voxel representation, which could still be able to encode sufficient geometric information, e.g., room layout, object's sizes and shapes, to infer the invisible areas of the scene with well structure-preserving details.

416, TITLE: Graph-Guided Architecture Search for Real-Time Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Graph-Guided_Architecture_Search_for_Real-Time_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Peiwen Lin, Peng Sun, Guangliang Cheng, Sirui Xie, Xi Li, Jianping Shi
HIGHLIGHT: In order to release researchers from these tedious mechanical trials, we propose a Graph-guided Architecture Search (GAS) pipeline to automatically search real-time semantic segmentation networks.

417, TITLE: Composing Good Shots by Exploiting Mutual Relations
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Composing_Good_Shots_by_Exploiting_Mutual_Relations_CVPR_2020_paper.html
AUTHORS: Debang Li, Junge Zhang, Kaiqi Huang, Ming-Hsuan Yang
HIGHLIGHT: Motivated by this, we propose a graph-based module with a gated feature update to model the relations between different candidates.

418, TITLE: Organ at Risk Segmentation for Head and Neck Cancer Using Stratified Learning and Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Organ_at_Risk_Segmentation_for_Head_and_Neck_Cancer_Using_CVPR_2020_paper.html
AUTHORS: Dazhou Guo, Dakai Jin, Zhuotun Zhu, Tsung-Ying Ho, Adam P. Harrison, Chun-Hung Chao, Jing Xiao, Le Lu
HIGHLIGHT: For such scenarios, insights can be gained from the stratification approaches seen in manual clinical OAR delineation. This is the goal of our work, where we introduce stratified organ at risk segmentation (SOARS), an approach that stratifies OARs into anchor, mid-level, and small & hard (S&H) categories.

419, TITLE: G2L-Net: Global to Local Network for Real-Time 6D Pose Estimation With Embedding Vector Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_G2L-Net_Global_to_Local_Network_for_Real-Time_6D_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, Ales Leonardis
HIGHLIGHT: In this paper, we propose a novel real-time 6D object pose estimation framework, named G2L-Net.

420, TITLE: Unsupervised Instance Segmentation in Microscopy Images via Panoptic Domain Adaptation and Task Re-Weighting
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Unsupervised_Instance_Segmentation_in_Microscopy_Images_via_Panoptic_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Dongnan Liu, Donghao Zhang, Yang Song, Fan Zhang, Lauren O'Donnell, Heng Huang, Mei Chen, Weidong Cai
HIGHLIGHT: In this work, we propose a Cycle Consistency Panoptic Domain Adaptive Mask R-CNN (CyC-PDAM) architecture for unsupervised nuclei segmentation in histopathology images, by learning from fluorescence microscopy images.

421, TITLE: Single-Stage Semantic Segmentation From Image Labels
http://openaccess.thecvf.com/content_CVPR_2020/html/Araslanov_Single-Stage_Semantic_Segmentation_From_Image_Labels_CVPR_2020_paper.html
AUTHORS: Nikita Araslanov, Stefan Roth
HIGHLIGHT: In this work, we first define three desirable properties of a weakly supervised method: local consistency, semantic fidelity, and completeness. Using these properties as guidelines, we then develop a segmentation-based network model and a self-supervised training scheme to train for semantic masks from image-level annotations in a single stage.

422, TITLE: Cascaded Human-Object Interaction Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Cascaded_Human-Object_Interaction_Recognition_CVPR_2020_paper.html
AUTHORS: Tianfei Zhou, Wenguan Wang, Siyuan Qi, Haibin Ling, Jianbing Shen
HIGHLIGHT: Considering the intrinsic complexity of the task, we introduce a cascade architecture for a multi-stage, coarse-to-fine HOI understanding.

423, TITLE: DuDoRNet: Learning a Dual-Domain Recurrent Network for Fast MRI Reconstruction With Deep T1 Prior
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_DuDoRNet_Learning_a_Dual-Domain_Recurrent_Network_for_Fast_MRI_Reconstruction_CVPR_2020_paper.html
AUTHORS: Bo Zhou, S. Kevin Zhou
HIGHLIGHT: In this work, we address the above two limitations by proposing a Dual Domain Recurrent Network (DuDoRNet) with deep T1 prior embedded to simultaneously recover k-space and images for accelerating the acquisition of MRI with a long imaging protocol.

424, TITLE: Learning Integral Objects With Intra-Class Discriminator for Weakly-Supervised Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_Learning_Integral_Objects_With_Intra-Class_Discriminator_for_Weakly-Supervised_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, Tieniu Tan
HIGHLIGHT: In this paper, we argue that the critical factor preventing to obtain the full object mask is the classification boundary mismatch problem in applying the CAM to WSSS.

425, TITLE: FPConv: Learning Local Flattening for Point Convolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_FPConv_Learning_Local_Flattening_for_Point_Convolution_CVPR_2020_paper.html
AUTHORS: Yiqun Lin, Zizheng Yan, Haibin Huang, Dong Du, Ligang Liu, Shuguang Cui, Xiaoguang Han
HIGHLIGHT: We introduce FPConv, a novel surface-style convolution operator designed for 3D point cloud analysis.

426, TITLE: Rotation Equivariant Graph Convolutional Network for Spherical Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Rotation_Equivariant_Graph_Convolutional_Network_for_Spherical_Image_Classification_CVPR_2020_paper.html
AUTHORS: Qin Yang, Chenglin Li, Wenrui Dai, Junni Zou, Guo-Jun Qi, Hongkai Xiong
HIGHLIGHT: In this paper, we generalize the grid-based CNNs to a non-Euclidean space by taking into account the geometry of spherical surfaces and propose a Spherical Graph Convolutional Network (SGCN) to encode rotation equivariant representations.

427, TITLE: FOAL: Fast Online Adaptive Learning for Cardiac Motion Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_FOAL_Fast_Online_Adaptive_Learning_for_Cardiac_Motion_Estimation_CVPR_2020_paper.html
AUTHORS: Hanchao Yu, Shanhui Sun, Haichao Yu, Xiao Chen, Honghui Shi, Thomas S. Huang, Terrence Chen
HIGHLIGHT: In this context, we proposed a novel fast online adaptive learning (FOAL) framework: an online gradient descent based optimizer that is optimized by a meta-learner.

428, TITLE: ScrabbleGAN: Semi-Supervised Varying Length Handwritten Text Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Fogel_ScrabbleGAN_Semi-Supervised_Varying_Length_Handwritten_Text_Generation_CVPR_2020_paper.html
AUTHORS: Sharon Fogel, Hadar Averbuch-Elor, Sarel Cohen, Shai Mazor, Roee Litman
HIGHLIGHT: We present ScrabbleGAN, a semi-supervised approach to synthesize handwritten text images that are versatile both in style and lexicon.

429, TITLE: Cross-Domain Semantic Segmentation via Domain-Invariant Interactive Relation Transfer
http://openaccess.thecvf.com/content_CVPR_2020/html/Lv_Cross-Domain_Semantic_Segmentation_via_Domain-Invariant_Interactive_Relation_Transfer_CVPR_2020_paper.html
AUTHORS: Fengmao Lv, Tao Liang, Xiang Chen, Guosheng Lin
HIGHLIGHT: In this paper, we propose a new domain adaptation approach, called Pivot Interaction Transfer (PIT).

430, TITLE: Inflated Episodic Memory With Region Self-Attention for Long-Tailed Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Inflated_Episodic_Memory_With_Region_Self-Attention_for_Long-Tailed_Visual_Recognition_CVPR_2020_paper.html
AUTHORS: Linchao Zhu, Yi Yang
HIGHLIGHT: To deal with the class imbalance problem, we introduce an Inflated Episodic Memory (IEM) for long-tailed visual recognition.

431, TITLE: Multimodal Future Localization and Emergence Prediction for Objects in Egocentric View With a Reachability Prior
http://openaccess.thecvf.com/content_CVPR_2020/html/Makansi_Multimodal_Future_Localization_and_Emergence_Prediction_for_Objects_in_Egocentric_CVPR_2020_paper.html
AUTHORS: Osama Makansi, Ozgun Cicek, Kevin Buchicchio, Thomas Brox
HIGHLIGHT: In this paper, we investigate the problem of anticipating future dynamics, particularly the future location of other vehicles and pedestrians, in the view of a moving vehicle.

432, TITLE: Structure Preserving Generative Cross-Domain Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Xia_Structure_Preserving_Generative_Cross-Domain_Learning_CVPR_2020_paper.html
AUTHORS: Haifeng Xia, Zhengming Ding
HIGHLIGHT: To this end, we develop a novel Generative cross-domain learning via Structure-Preserving (GSP), which attempts to transform target data into the source domain in order to take advantage of source supervision.

433, TITLE: Reverse Perspective Network for Perspective-Aware Object Counting
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Reverse_Perspective_Network_for_Perspective-Aware_Object_Counting_CVPR_2020_paper.html
AUTHORS: Yifan Yang, Guorong Li, Zhe Wu, Li Su, Qingming Huang, Nicu Sebe
HIGHLIGHT: We propose a reverse perspective network to solve the scale variations of input images, instead of generating perspective maps to smooth final outputs.

434, TITLE: Multi-Path Region Mining for Weakly Supervised 3D Semantic Segmentation on Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Multi-Path_Region_Mining_for_Weakly_Supervised_3D_Semantic_Segmentation_on_CVPR_2020_paper.html
AUTHORS: Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Tzu-Yi Hung, Lihua Xie
HIGHLIGHT: In this paper, we propose a weakly supervised approach to predict point-level results using weak labels on 3D point clouds.

435, TITLE: Reliable Weighted Optimal Transport for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Reliable_Weighted_Optimal_Transport_for_Unsupervised_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Renjun Xu, Pelen Liu, Liyan Wang, Chao Chen, Jindong Wang
HIGHLIGHT: In this paper, we present Reliable Weighted Optimal Transport (RWOT) for unsupervised domain adaptation, including novel Shrinking Subspace Reliability (SSR) and weighted optimal transport strategy.

436, TITLE: ImVoteNet: Boosting 3D Object Detection in Point Clouds With Image Votes
http://openaccess.thecvf.com/content_CVPR_2020/html/Qi_ImVoteNet_Boosting_3D_Object_Detection_in_Point_Clouds_With_Image_CVPR_2020_paper.html
AUTHORS: Charles R. Qi, Xinlei Chen, Or Litany, Leonidas J. Guibas
HIGHLIGHT: In this work, we build on top of VoteNet and propose a 3D detection architecture called ImVoteNet specialized for RGB-D scenes.

437, TITLE: Understanding Road Layout From Videos as a Whole
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Understanding_Road_Layout_From_Videos_as_a_Whole_CVPR_2020_paper.html
AUTHORS: Buyu Liu, Bingbing Zhuang, Samuel Schulter, Pan Ji, Manmohan Chandraker
HIGHLIGHT: In this paper, we address the problem of inferring the layout of complex road scenes from video sequences.

438, TITLE: Bi-Directional Relationship Inferring Network for Referring Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Bi-Directional_Relationship_Inferring_Network_for_Referring_Image_Segmentation_CVPR_2020_paper.html
AUTHORS: Zhiwei Hu, Guang Feng, Jiayu Sun, Lihe Zhang, Huchuan Lu
HIGHLIGHT: In this work, we propose a bi-directional relationship inferring network (BRINet) to model the dependencies of cross-modal information.

439, TITLE: Perspective Plane Program Induction From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Perspective_Plane_Program_Induction_From_a_Single_Image_CVPR_2020_paper.html
AUTHORS: Yikai Li, Jiayuan Mao, Xiuming Zhang, William T. Freeman, Joshua B. Tenenbaum, Jiajun Wu

HIGHLIGHT: We formulate this problem as jointly finding the camera pose and scene structure that best describe the input image.

440, TITLE: DeepFLASH: An Efficient Network for Learning-Based Medical Image Registration
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_DeepFLASH_An_Efficient_Network_for_Learning-Based_Medical_Image_Registration_CVPR_2020_paper.html
AUTHORS: Jian Wang, Miaomiao Zhang
HIGHLIGHT: This paper presents DeepFLASH, a novel network with efficient training and inference for learning-based medical image registration.

441, TITLE: Semi-Supervised Learning for Few-Shot Image-to-Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Semi-Supervised_Learning_for_Few-Shot_Image-to-Image_Translation_CVPR_2020_paper.html
AUTHORS: Yaxing Wang, Salman Khan, Abel Gonzalez-Garcia, Joost van de Weijer, Fahad Shahbaz Khan
HIGHLIGHT: In this work, we go one step further and reduce the amount of required labeled data also from the source domain during training.

442, TITLE: Semantic Correspondence as an Optimal Transport Problem
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Semantic_Correspondence_as_an_Optimal_Transport_Problem_CVPR_2020_paper.html
AUTHORS: Yanbin Liu, Linchao Zhu, Makoto Yamada, Yi Yang
HIGHLIGHT: The whole procedure is combined into a unified optimal transport algorithm by converting the maximization problem to the optimal transport formulation and incorporating the staircase weights into optimal transport algorithm to act as empirical distributions.

443, TITLE: How Much Time Do You Have? Modeling Multi-Duration Saliency
http://openaccess.thecvf.com/content_CVPR_2020/html/Fosco_How_Much_Time_Do_You_Have_Modeling_Multi-Duration_Saliency_CVPR_2020_paper.html
AUTHORS: Camilo Fosco, Anelise Newman, Pat Sukhum, Yun Bin Zhang, Nanxuan Zhao, Aude Oliva, Zoya Bylinskii
HIGHLIGHT: In this paper we propose to capture gaze as a series of snapshots, by generating population-level saliency heatmaps for multiple viewing durations.
We collect the CodeCharts1K dataset, which contains multiple distinct heatmaps per image corresponding to 0.5, 3, and 5 seconds of free-viewing.

444, TITLE: Fine-Grained Generalized Zero-Shot Learning via Dense Attribute-Based Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Huynh_Fine-Grained_Generalized_Zero-Shot_Learning_via_Dense_Attribute-Based_Attention_CVPR_2020_paper.html
AUTHORS: Dat Huynh, Ehsan Elhamifar
HIGHLIGHT: Instead of aligning a global feature vector of an image with its associated class semantic vector, we propose an attribute embedding technique that aligns each attribute-based feature with its attribute semantic vector.

445, TITLE: Online Depth Learning Against Forgetting in Monocular Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Online_Depth_Learning_Against_Forgetting_in_Monocular_Videos_CVPR_2020_paper.html
AUTHORS: Zhenyu Zhang, Stephane Lathuiliere, Elisa Ricci, Nicu Sebe, Yan Yan, Jian Yang
HIGHLIGHT: Specifically, to adapt temporal-continuous depth patterns in videos, we introduce a novel meta-learning approach to learn adapter modules by combining online adaptation process into the learning objective.

446, TITLE: Few-Shot Learning of Part-Specific Probability Space for 3D Shape Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Few-Shot_Learning_of_Part-Specific_Probability_Space_for_3D_Shape_Segmentation_CVPR_2020_paper.html
AUTHORS: Lingjing Wang, Xiang Li, Yi Fang
HIGHLIGHT: In comparison, we propose a novel 3D shape segmentation method that requires few labeled data for training.

447, TITLE: Pattern-Structure Diffusion for Multi-Task Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Pattern-Structure_Diffusion_for_Multi-Task_Learning_CVPR_2020_paper.html
AUTHORS: Ling Zhou, Zhen Cui, Chunyan Xu, Zhenyu Zhang, Chaoqun Wang, Tong Zhang, Jian Yang
HIGHLIGHT: Inspired by the observation that pattern structures high-frequently recur within intra-task also across tasks, we propose a pattern-structure diffusion (PSD) framework to mine and propagate task-specific and task-across pattern structures in the task-level space for joint depth estimation, segmentation and surface normal prediction.

448, TITLE: Training Noise-Robust Deep Neural Networks via Meta-Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Training_Noise-Robust_Deep_Neural_Networks_via_Meta-Learning_CVPR_2020_paper.html
AUTHORS: Zhen Wang, Guosheng Hu, Qinghua Hu
HIGHLIGHT: In this work, we propose a new loss correction approach, named as Meta Loss Correction (MLC), to directly learn T from data via the meta-learning framework.

449, TITLE: Fusion-Aware Point Convolution for Online Semantic 3D Scene Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Fusion-Aware_Point_Convolution_for_Online_Semantic_3D_Scene_Segmentation_CVPR_2020_paper.html
AUTHORS: Jiazhao Zhang, Chenyang Zhu, Lintao Zheng, Kai Xu
HIGHLIGHT: We propose a novel fusion-aware 3D point convolution which operates directly on the geometric surface being reconstructed and exploits effectively the inter-frame correlation for high-quality 3D feature learning.

450, TITLE: Universal Source-Free Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kundu_Universal_Source-Free_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Jogendra Nath Kundu, Naveen Venkat, Rahul M V, R. Venkatesh Babu
HIGHLIGHT: Devoid of such impractical assumptions, we propose a novel two-stage learning process.

451, TITLE: Exploring Spatial-Temporal Multi-Frequency Analysis for High-Fidelity and Temporal-Consistency Video Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Jin_Exploring_Spatial-Temporal_Multi-Frequency_Analysis_for_High-Fidelity_and_Temporal-Consistency_Video_Prediction_CVPR_2020_paper.html
AUTHORS: Beibei Jin, Yu Hu, Qiankun Tang, Jingyu Niu, Zhiping Shi, Yinhe Han, Xiaowei Li
HIGHLIGHT: Inspired by the frequency band decomposition characteristic of Human Vision System (HVS), we propose a video prediction network based on multi-level wavelet analysis to uniformly deal with spatial and temporal information.

452, TITLE: Varicolored Image De-Hazing
http://openaccess.thecvf.com/content_CVPR_2020/html/Dudhane_Varicolored_Image_De-Hazing_CVPR_2020_paper.html
AUTHORS: Akshay Dudhane, Kuldeep M. Biradar, Prashant W. Patil, Praful Hambarde, Subrahmanyam Murala
HIGHLIGHT: In this paper, we propose a varicolored end-to-end image de-hazing network which restores the color balance in a given varicolored hazy image and recovers the haze-free image.

453, TITLE: SpSequenceNet: Semantic Segmentation Network on 4D Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_SpSequenceNet_Semantic_Segmentation_Network_on_4D_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Hanyu Shi, Guosheng Lin, Hao Wang, Tzu-Yi Hung, Zhenhua Wang
HIGHLIGHT: In this paper, we propose SpSequenceNet to address this problem.

454, TITLE: Separating Particulate Matter From a Single Microscopic Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Sandhan_Separating_Particulate_Matter_From_a_Single_Microscopic_Image_CVPR_2020_paper.html
AUTHORS: Tushar Sandhan, Jin Young Choi
HIGHLIGHT: In this work, we thoroughly analyze the physical properties of PM, microscope and their inevitable interaction; and propose an optimization scheme, which removes the PM from a high-resolution microscopic image within a few seconds.

455, TITLE: Adaptive Dilated Network With Self-Correction Supervision for Counting
http://openaccess.thecvf.com/content_CVPR_2020/html/Bai_Adaptive_Dilated_Network_With_Self-Correction_Supervision_for_Counting_CVPR_2020_paper.html
AUTHORS: Shuai Bai, Zhiqun He, Yu Qiao, Hanzhe Hu, Wei Wu, Junjie Yan
HIGHLIGHT: In this paper, we propose an adaptive dilated convolution and a novel supervised learning framework named self-correction (SC) supervision.

456, TITLE: PointPainting: Sequential Fusion for 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Vora_PointPainting_Sequential_Fusion_for_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Sourabh Vora, Alex H. Lang, Bassam Helou, Oscar Beijbom
HIGHLIGHT: In this work, we propose PointPainting: a sequential fusion method to fill this gap.

457, TITLE: Rethinking Zero-Shot Video Classification: End-to-End Training for Realistic Applications
http://openaccess.thecvf.com/content_CVPR_2020/html/Brattoli_Rethinking_Zero-Shot_Video_Classification_End-to-End_Training_for_Realistic_Applications_CVPR_2020_paper.html
AUTHORS: Biagio Brattoli, Joseph Tighe, Fedor Zhdanov, Pietro Perona, Krzysztof Chalupka
HIGHLIGHT: We propose the first end-to-end algorithm for ZSL in video classification.

458, TITLE: Learning to Select Base Classes for Few-Shot Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Learning_to_Select_Base_Classes_for_Few-Shot_Classification_CVPR_2020_paper.html
AUTHORS: Linjun Zhou, Peng Cui, Xu Jia, Shiqiang Yang, Qi Tian
HIGHLIGHT: In this paper, we utilize a simple yet effective measure, the Similarity Ratio, as an indicator for the generalization performance of a few-shot model. We then formulate the base class selection problem as a submodular optimization problem over Similarity Ratio.

459, TITLE: CONSAC: Robust Multi-Model Fitting by Conditional Sample Consensus
http://openaccess.thecvf.com/content_CVPR_2020/html/Kluger_CONSAC_Robust_Multi-Model_Fitting_by_Conditional_Sample_Consensus_CVPR_2020_paper.html
AUTHORS: Florian Kluger, Eric Brachmann, Hanno Ackermann, Carsten Rother, Michael Ying Yang, Bodo Rosenhahn
HIGHLIGHT: We present a robust estimator for fitting multiple parametric models of the same form to noisy measurements.

460, TITLE: Fast Symmetric Diffeomorphic Image Registration with Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Mok_Fast_Symmetric_Diffeomorphic_Image_Registration_with_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Tony C.W. Mok, Albert C.S. Chung
HIGHLIGHT: In this paper, we present a novel, efficient unsupervised symmetric image registration method which maximizes the similarity between images within the space of diffeomorphic maps and estimates both forward and inverse transformations simultaneously.

461, TITLE: Distilled Semantics for Comprehensive Scene Understanding from Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Tosi_Distilled_Semantics_for_Comprehensive_Scene_Understanding_from_Videos_CVPR_2020_paper.html
AUTHORS: Fabio Tosi, Filippo Aleotti, Pierluigi Zama Ramirez, Matteo Poggi, Samuele Salti, Luigi Di Stefano, Stefano Mattoccia
HIGHLIGHT: In this paper, we take an additional step toward holistic scene understanding with monocular cameras by learning depth and motion alongside with semantics, with supervision for the latter provided by a pre-trained network distilling proxy ground truth images.

462, TITLE: Modeling Biological Immunity to Adversarial Examples
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Modeling_Biological_Immunity_to_Adversarial_Examples_CVPR_2020_paper.html
AUTHORS: Edward Kim, Jocelyn Rego, Yijing Watkins, Garrett T. Kenyon
HIGHLIGHT: In this work, we explored this gap through the lens of biology and neuroscience in order to understand the robustness exhibited in human perception.

463, TITLE: DOA-GAN: Dual-Order Attentive Generative Adversarial Network for Image Copy-Move Forgery Detection and Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Islam_DOA-GAN_Dual-Order_Attentive_Generative_Adversarial_Network_for_Image_Copy-Move_Forgery_CVPR_2020_paper.html
AUTHORS: Ashraful Islam, Chengjiang Long, Arslan Basharat, Anthony Hoogs
HIGHLIGHT: In this paper, we propose a Generative Adversarial Network with a dual-order attention model to detect and localize copy-move forgeries.

464, TITLE: Correspondence-Free Material Reconstruction using Sparse Surface Constraints
http://openaccess.thecvf.com/content_CVPR_2020/html/Weiss_Correspondence-Free_Material_Reconstruction_using_Sparse_Surface_Constraints_CVPR_2020_paper.html
AUTHORS: Sebastian Weiss, Robert Maier, Daniel Cremers, Rudiger Westermann, Nils Thuerey
HIGHLIGHT: We present a method to infer physical material parameters, and even external boundaries, from the scanned motion of a homogeneous deformable object via the solution of an inverse problem.

465, TITLE: Augmenting Colonoscopy Using Extended and Directional CycleGAN for Lossy Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Mathew_Augmenting_Colonoscopy_Using_Extended_and_Directional_CycleGAN_for_Lossy_Image_CVPR_2020_paper.html

AUTHORS: Shawn Mathew, Saad Nadeem, Sruti Kumari, Arie Kaufman
HIGHLIGHT: In this paper, we present a deep learning framework, Extended and Directional CycleGAN, for lossy unpaired image-to-image translation between OC and VC to augment OC video sequences with scale-consistent depth information from VC and VC with patient-specific textures, color and specular highlights from OC (e.g. for realistic polyp synthesis).

466, TITLE: Attention Scaling for Crowd Counting
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Attention_Scaling_for_Crowd_Counting_CVPR_2020_paper.html
AUTHORS: Xiaoheng Jiang, Li Zhang, Mingliang Xu, Tianzhu Zhang, Pei Lv, Bing Zhou, Xin Yang, Yanwei Pang
HIGHLIGHT: To overcome this problem, we propose an approach to alleviate the counting performance differences in different regions.

467, TITLE: Shape Reconstruction by Learning Differentiable Surface Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Bednarik_Shape_Reconstruction_by_Learning_Differentiable_Surface_Representations_CVPR_2020_paper.html
AUTHORS: Jan Bednarik, Shaifali Parashar, Erhan Gundogdu, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: In this paper, we show that we can exploit the inherent differentiability of deep networks to leverage differential surface properties during training so as to prevent patch collapse and strongly reduce patch overlap.

468, TITLE: A Spatiotemporal Volumetric Interpolation Network for 4D Dynamic Medical Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_A_Spatiotemporal_Volumetric_Interpolation_Network_for_4D_Dynamic_Medical_Image_CVPR_2020_paper.html
AUTHORS: Yuyu Guo, Lei Bi, Euijoon Ahn, Dagan Feng, Qian Wang, Jinman Kim
HIGHLIGHT: In this paper, we present a spatiotemporal volumetric interpolation network (SVIN) designed for 4D dynamic medical images.

469, TITLE: Attention-Based Context Aware Reasoning for Situation Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Cooray_Attention-Based_Context_Aware_Reasoning_for_Situation_Recognition_CVPR_2020_paper.html
AUTHORS: Thilini Cooray, Ngai-Man Cheung, Wei Lu
HIGHLIGHT: Inspired by the success achieved by query-based visual reasoning (e.g., Visual Question Answering), we propose to address semantic role prediction as a query-based visual reasoning problem.

470, TITLE: PatchVAE: Learning Local Latent Codes for Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Gupta_PatchVAE_Learning_Local_Latent_Codes_for_Recognition_CVPR_2020_paper.html
AUTHORS: Kamal Gupta, Saurabh Singh, Abhinav Shrivastava
HIGHLIGHT: Drawing inspiration from the mid-level representation discovery work, we propose PatchVAE, that reasons about images at patch level.

471, TITLE: Self-Supervised Monocular Trained Depth Estimation Using Self-Attention and Discrete Disparity Volume
http://openaccess.thecvf.com/content_CVPR_2020/html/Johnston_Self-Supervised_Monocular_Trained_Depth_Estimation_Using_Self-Attention_and_Discrete_Disparity_CVPR_2020_paper.html
AUTHORS: Adrian Johnston, Gustavo Carneiro
HIGHLIGHT: In this paper, we propose two new ideas to improve self-supervised monocular trained depth estimation: 1) self-attention, and 2) discrete disparity prediction.

472, TITLE: STAViS: Spatio-Temporal AudioVisual Saliency Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Tsiami_STAViS_Spatio-Temporal_AudioVisual_Saliency_Network_CVPR_2020_paper.html
AUTHORS: Antigoni Tsiami, Petros Koutras, Petros Maragos
HIGHLIGHT: We introduce STAViS, a spatio-temporal audiovisual saliency network that combines spatio-temporal visual and auditory information in order to efficiently address the problem of saliency estimation in videos.

473, TITLE: More Grounded Image Captioning by Distilling Image-Text Matching Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_More_Grounded_Image_Captioning_by_Distilling_Image-Text_Matching_Model_CVPR_2020_paper.html
AUTHORS: Yuanen Zhou, Meng Wang, Daqing Liu, Zhenzhen Hu, Hanwang Zhang
HIGHLIGHT: To this end, we propose a Part-of-Speech (POS) enhanced image-text matching model (SCAN): POS-SCAN, as the effective knowledge distillation for more grounded image captioning.

474, TITLE: DUNIT: Detection-Based Unsupervised Image-to-Image Translation

http://openaccess.thecvf.com/content_CVPR_2020/html/Bhattacharjee_DUNIT_Detection-Based_Unsupervised_Image-to-Image_Translation_CVPR_2020_paper.html
AUTHORS: Deblina Bhattacharjee, Seungryong Kim, Guillaume Vizier, Mathieu Salzmann
HIGHLIGHT: In this paper, we introduce a Detection-based Unsupervised Image-to-image Translation (DUNIT) approach that explicitly accounts for the object instances in the translation process.

475, TITLE: Learning to Observe: Approximating Human Perceptual Thresholds for Detection of Suprathreshold Image Transformations
http://openaccess.thecvf.com/content_CVPR_2020/html/Dolhasz_Learning_to_Observe_Approximating_Human_Perceptual_Thresholds_for_Detection_of_CVPR_2020_paper.html
AUTHORS: Alan Dolhasz, Carlo Harvey, Ian Williams
HIGHLIGHT: In this paper, we propose to directly approximate the perceptual function performed by human observers completing a visual detection task.

476, TITLE: Show, Edit and Tell: A Framework for Editing Image Captions
http://openaccess.thecvf.com/content_CVPR_2020/html/Sammani_Show_Edit_and_Tell_A_Framework_for_Editing_Image_Captions_CVPR_2020_paper.html
AUTHORS: Fawaz Sammani, Luke Melas-Kyriazi
HIGHLIGHT: This paper proposes a novel approach to image captioning based on iterative adaptive refinement of an existing caption.

477, TITLE: Structure Boundary Preserving Segmentation for Medical Image With Ambiguous Boundary
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Structure_Boundary_Preserving_Segmentation_for_Medical_Image_With_Ambiguous_Boundary_CVPR_2020_paper.html
AUTHORS: Hong Joo Lee, Jung Uk Kim, Sangmin Lee, Hak Gu Kim, Yong Man Ro
HIGHLIGHT: In this paper, we propose a novel image segmentation method to tackle two critical problems of medical image, which are (i) ambiguity of structure boundary in the medical image domain and (ii) uncertainty of the segmented region without specialized domain knowledge.

478, TITLE: Predicting Cognitive Declines Using Longitudinally Enriched Representations for Imaging Biomarkers
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Predicting_Cognitive_Declines_Using_Longitudinally_Enriched_Representations_for_Imaging_Biomarkers_CVPR_2020_paper.html
AUTHORS: Lyujian Lu, Hua Wang, Saad Elbeleidy, Feiping Nie
HIGHLIGHT: To tackle this problem, in this paper we propose a novel formulation to learn an enriched representation for imaging biomarkers that can simultaneously capture both the information conveyed by baseline neuroimaging records and that by progressive variations of varied counts of available follow-up records over time.

479, TITLE: Predicting Lymph Node Metastasis Using Histopathological Images Based on Multiple Instance Learning With Deep Graph Convolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Predicting_Lymph_Node_Metastasis_Using_Histopathological_Images_Based_on_Multiple_CVPR_2020_paper.html
AUTHORS: Yu Zhao, Fan Yang, Yuqi Fang, Hailing Liu, Niyun Zhou, Jun Zhang, Jiarui Sun, Sen Yang, Bjoern Menze, Xinjuan Fan, Jianhua Yao
HIGHLIGHT: In this paper, we propose a multiple instance learning method based on deep graph convolutional network and feature selection (FS-GCN-MIL) for histopathological image classification.

480, TITLE: Extremely Dense Point Correspondences Using a Learned Feature Descriptor
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Extremely_Dense_Point_Correspondences_Using_a_Learned_Feature_Descriptor_CVPR_2020_paper.html
AUTHORS: Xingtong Liu, Yiping Zheng, Benjamin Killeen, Masaru Ishii, Gregory D. Hager, Russell H. Taylor, Mathias Unberath
HIGHLIGHT: In this work, we present an effective self-supervised training scheme and novel loss design for dense descriptor learning.

481, TITLE: Local Deep Implicit Functions for 3D Shape
http://openaccess.thecvf.com/content_CVPR_2020/html/Genova_Local_Deep_Implicit_Functions_for_3D_Shape_CVPR_2020_paper.html
AUTHORS: Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, Thomas Funkhouser
HIGHLIGHT: Towards this end, we introduce Local Deep Implicit Functions (LDIF), a 3D shape representation that decomposes space into a structured set of learned implicit functions.

482, TITLE: PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation

http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_PointGroup_Dual-Set_Point_Grouping_for_3D_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, Jiaya Jia
HIGHLIGHT: In this paper, we present PointGroup, a new end-to-end bottom-up architecture, specifically focused on better grouping the points by exploring the void space between objects.

483, TITLE: Cost Volume Pyramid Based Depth Inference for Multi-View Stereo
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Cost_Volume_Pyramid_Based_Depth_Inference_for_Multi-View_Stereo_CVPR_2020_paper.html
AUTHORS: Jiayu Yang, Wei Mao, Jose M. Alvarez, Miaomiao Liu
HIGHLIGHT: We propose a cost volume-based neural network for depth inference from multi-view images.

484, TITLE: RoutedFusion: Learning Real-Time Depth Map Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Weder_RoutedFusion_Learning_Real-Time_Depth_Map_Fusion_CVPR_2020_paper.html
AUTHORS: Silvan Weder, Johannes Schonberger, Marc Pollefeys, Martin R. Oswald
HIGHLIGHT: To this end, we present a novel real-time capable machine learning-based method for depth map fusion.

485, TITLE: VOLDOR: Visual Odometry From Log-Logistic Dense Optical Flow Residuals
http://openaccess.thecvf.com/content_CVPR_2020/html/Min_VOLDOR_Visual_Odometry_From_Log-Logistic_Dense_Optical_Flow_Residuals_CVPR_2020_paper.html
AUTHORS: Zhixiang Min, Yiding Yang, Enrique Dunn
HIGHLIGHT: We propose a dense indirect visual odometry method taking as input externally estimated optical flow fields instead of hand-crafted feature correspondences.

486, TITLE: Learning to Optimize Non-Rigid Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Learning_to_Optimize_Non-Rigid_Tracking_CVPR_2020_paper.html
AUTHORS: Yang Li, Aljaz Bozic, Tianwei Zhang, Yanli Ji, Tatsuya Harada, Matthias Niessner
HIGHLIGHT: In this paper, we employ learnable optimizations to improve tracking robustness and speed up solver convergence.

487, TITLE: KFNet: Learning Temporal Camera Relocalization Using Kalman Filtering
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_KFNet_Learning_Temporal_Camera_Relocalization_Using_Kalman_Filtering_CVPR_2020_paper.html
AUTHORS: Lei Zhou, Zixin Luo, Tianwei Shen, Jiahui Zhang, Mingmin Zhen, Yao Yao, Tian Fang, Long Quan
HIGHLIGHT: In this work, we improve the temporal relocalization method by using a network architecture that incorporates Kalman filtering (KFNet) for online camera relocalization.

488, TITLE: Information-Driven Direct RGB-D Odometry
http://openaccess.thecvf.com/content_CVPR_2020/html/Fontan_Information-Driven_Direct_RGB-D_Odometry_CVPR_2020_paper.html
AUTHORS: Alejandro Fontan, Javier Civera, Rudolph Triebel
HIGHLIGHT: This paper presents an information-theoretic approach to point selection in direct RGB-D odometry.

489, TITLE: SuperGlue: Learning Feature Matching With Graph Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Sarlin_SuperGlue_Learning_Feature_Matching_With_Graph_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, Andrew Rabinovich
HIGHLIGHT: This paper introduces SuperGlue, a neural network that matches two sets of local features by jointly finding correspondences and rejecting non-matchable points.

490, TITLE: Reinforced Feature Points: Optimizing Feature Detection and Description for a High-Level Task
http://openaccess.thecvf.com/content_CVPR_2020/html/Bhowmik_Reinforced_Feature_Points_Optimizing_Feature_Detection_and_Description_for_a_CVPR_2020_paper.html
AUTHORS: Aritra Bhowmik, Stefan Gumhold, Carsten Rother, Eric Brachmann
HIGHLIGHT: We propose a new training methodology which embeds the feature detector in a complete vision pipeline, and where the learnable parameters are trained in an end-to-end fashion.

491, TITLE: ReDA:Reinforced Differentiable Attribute for 3D Face Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_ReDAReinforced_Differentiable_Attribute_for_3D_Face_Reconstruction_CVPR_2020_paper.html

AUTHORS: Wenbin Zhu, HsiangTao Wu, Zeyu Chen, Noranart Vesdapunt, Baoyuan Wang
HIGHLIGHT: To further reduce the ambiguities, we present a novel framework called "Reinforced Differentiable Attributes" ("ReDA") which is more general and effective than previous Differentiable Rendering ("DR").

492, TITLE: EventCap: Monocular 3D Capture of High-Speed Human Motions Using an Event Camera
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_EventCap_Monocular_3D_Capture_of_High-Speed_Human_Motions_Using_an_CVPR_2020_paper.html
AUTHORS: Lan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, Christian Theobalt
HIGHLIGHT: In this paper, we propose EventCap -- the first approach for 3D capturing of high-speed human motions using a single event camera.

493, TITLE: Cross-Modal Deep Face Normals With Deactivable Skip Connections
http://openaccess.thecvf.com/content_CVPR_2020/html/Abrevaya_Cross-Modal_Deep_Face_Normals_With_Deactivable_Skip_Connections_CVPR_2020_paper.html
AUTHORS: Victoria Fernandez Abrevaya, Adnane Boukhayma, Philip H.S. Torr, Edmond Boyer
HIGHLIGHT: We present an approach for estimating surface normals from in-the-wild color images of faces.

494, TITLE: Weakly-Supervised Mesh-Convolutional Hand Reconstruction in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Kulon_Weakly-Supervised_Mesh-Convolutional_Hand_Reconstruction_in_the_Wild_CVPR_2020_paper.html
AUTHORS: Dominik Kulon, Riza Alp Guler, Iasonas Kokkinos, Michael M. Bronstein, Stefanos Zafeiriou
HIGHLIGHT: We introduce a simple and effective network architecture for monocular 3D hand pose estimation consisting of an image encoder followed by a mesh convolutional decoder that is trained through a direct 3D hand mesh reconstruction loss.

495, TITLE: Face X-Ray for More General Face Forgery Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Face_X-Ray_for_More_General_Face_Forgery_Detection_CVPR_2020_paper.html
AUTHORS: Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, Baining Guo
HIGHLIGHT: In this paper we propose a novel image representation called face X-ray for detecting forgery in face images.

496, TITLE: A Morphable Face Albedo Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Smith_A_Morphable_Face_Albedo_Model_CVPR_2020_paper.html
AUTHORS: William A. P. Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua B. Tenenbaum, Bernhard Egger
HIGHLIGHT: In this paper, we bring together two divergent strands of research: photometric face capture and statistical 3D face appearance modelling.

497, TITLE: Cascade EF-GAN: Progressive Facial Expression Editing With Local Focuses
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Cascade_EF-GAN_Progressive_Facial_Expression_Editing_With_Local_Focuses_CVPR_2020_paper.html
AUTHORS: Rongliang Wu, Gongjie Zhang, Shijian Lu, Tao Chen
HIGHLIGHT: To address these limitations, we propose Cascade Expression Focal GAN (Cascade EF-GAN), a novel network that performs progressive facial expression editing with local expression focuses.

498, TITLE: GanHand: Predicting Human Grasp Affordances in Multi-Object Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Corona_GanHand_Predicting_Human_Grasp_Affordances_in_Multi-Object_Scenes_CVPR_2020_paper.html
AUTHORS: Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer, Gregory Rogez
HIGHLIGHT: To this end, we introduce a generative model that jointly reasons in all these levels and 1) regresses the 3D shape and pose of the objects in the scene; 2) estimates the grasp types; and 3) refines the 51-DoF of a 3D hand model that minimize a graspability loss.

499, TITLE: Deep Spatial Gradient and Temporal Depth Learning for Face Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Deep_Spatial_Gradient_and_Temporal_Depth_Learning_for_Face_Anti-Spoofing_CVPR_2020_paper.html
AUTHORS: Zezheng Wang, Zitong Yu, Chenxu Zhao, Xiangyu Zhu, Yunxiao Qin, Qiusheng Zhou, Feng Zhou, Zhen Lei
HIGHLIGHT: In contrast, we design a new approach to detect presentation attacks from multiple frames based on two insights: 1) detailed discriminative clues (e.g., spatial gradient magnitude) between living and spoofing face may be discarded through stacked vanilla convolutions, and 2) the dynamics of 3D moving faces provide important clues in detecting the spoofing faces.

500, TITLE: DeepCap: Monocular Human Performance Capture Using Weak Supervision

http://openaccess.thecvf.com/content_CVPR_2020/html/Habermann_DeepCap_Monocular_Human_Performance_Capture_Using_Weak_Supervision_CVPR_2020_paper.html
AUTHORS: Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, Christian Theobalt
HIGHLIGHT: We propose a novel deep learning approach for monocular dense human performance capture.

501, TITLE: Attention Mechanism Exploits Temporal Contexts: Real-Time 3D Human Pose Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Attention_Mechanism_Exploits_Temporal_Contexts_Real-Time_3D_Human_Pose_Reconstruction_CVPR_2020_paper.html
AUTHORS: Ruixu Liu, Ju Shen, He Wang, Chen Chen, Sen-ching Cheung, Vijayan Asari
HIGHLIGHT: We propose a novel attention-based framework for 3D human pose estimation from a monocular video.

502, TITLE: Advancing High Fidelity Identity Swapping for Forgery Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Advancing_High_Fidelity_Identity_Swapping_for_Forgery_Detection_CVPR_2020_paper.html
AUTHORS: Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen
HIGHLIGHT: In this work, we study various existing benchmarks for deepfake detection researches.

503, TITLE: Controllable Person Image Synthesis With Attribute-Decomposed GAN
http://openaccess.thecvf.com/content_CVPR_2020/html/Men_Controllable_Person_Image_Synthesis_With_Attribute-Decomposed_GAN_CVPR_2020_paper.html
AUTHORS: Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, Zhouhui Lian
HIGHLIGHT: This paper introduces the Attribute-Decomposed GAN, a novel generative model for controllable person image synthesis, which can produce realistic person images with desired human attributes (e.g., pose, head, upper clothes and pants) provided in various source inputs.

504, TITLE: Attentive Normalization for Conditional Image Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Attentive_Normalization_for_Conditional_Image_Generation_CVPR_2020_paper.html
AUTHORS: Yi Wang, Ying-Cong Chen, Xiangyu Zhang, Jian Sun, Jiaya Jia
HIGHLIGHT: In this paper, we characterize long-range dependence with attentive normalization (AN), which is an extension to traditional instance normalization.

505, TITLE: SEAN: Image Synthesis With Semantic Region-Adaptive Normalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_SEAN_Image_Synthesis_With_Semantic_Region-Adaptive_Normalization_CVPR_2020_paper.html
AUTHORS: Peihao Zhu, Rameen Abdal, Yipeng Qin, Peter Wonka
HIGHLIGHT: We propose semantic region-adaptive normalization (SEAN), a simple but effective building block for Generative Adversarial Networks conditioned on segmentation masks that describe the semantic regions in the desired output image.

506, TITLE: Blurry Video Frame Interpolation
http://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Blurry_Video_Frame_Interpolation_CVPR_2020_paper.html
AUTHORS: Wang Shen, Wenbo Bao, Guangtao Zhai, Li Chen, Xiongkuo Min, Zhiyong Gao
HIGHLIGHT: In this paper, we propose a blurry video frame interpolation method to reduce motion blur and up-convert frame rate simultaneously.

507, TITLE: Learning Physics-Guided Face Relighting Under Directional Light
http://openaccess.thecvf.com/content_CVPR_2020/html/Nestmeyer_Learning_Physics-Guided_Face_Relighting_Under_Directional_Light_CVPR_2020_paper.html
AUTHORS: Thomas Nestmeyer, Jean-Francois Lalonde, Iain Matthews, Andreas Lehrmann
HIGHLIGHT: We investigate end-to-end deep learning architectures that both de-light and relight an image of a human face.

508, TITLE: Disentangled Image Generation Through Structured Noise Injection
http://openaccess.thecvf.com/content_CVPR_2020/html/Alharbi_Disentangled_Image_Generation_Through_Structured_Noise_Injection_CVPR_2020_paper.html
AUTHORS: Yazeed Alharbi, Peter Wonka
HIGHLIGHT: Instead of traditional approaches, we propose feeding multiple noise codes through separate fully-connected layers respectively.

509, TITLE: Cross-Domain Correspondence Learning for Exemplar-Based Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Cross-Domain_Correspondence_Learning_for_Exemplar-Based_Image_Translation_CVPR_2020_paper.html

AUTHORS: Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, Fang Wen
HIGHLIGHT: We present a general framework for exemplar-based image translation, which synthesizes a photo-realistic image from the input in a distinct domain (e.g., semantic segmentation mask, or edge map, or pose keypoints), given an exemplar image.

510, TITLE: Disentangled and Controllable Face Image Generation via 3D Imitative-Contrastive Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Deng_Disentangled_and_Controllable_Face_Image_Generation_via_3D_Imitative-Contrastive_Learning_CVPR_2020_paper.html
AUTHORS: Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, Xin Tong
HIGHLIGHT: We propose an approach for face image generation of virtual people with disentangled, precisely-controllable latent representations for identity of non-existing people, expression, pose, and illumination.

511, TITLE: Single Image Reflection Removal With Physically-Based Training Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Single_Image_Reflection_Removal_With_Physically-Based_Training_Images_CVPR_2020_paper.html
AUTHORS: Soomin Kim, Yuchi Huo, Sung-Eui Yoon
HIGHLIGHT: In this paper, physically based rendering is used for faithfully synthesizing the required training images, and a corresponding network structure and loss term are proposed.

512, TITLE: SketchyCOCO: Image Generation From Freehand Scene Sketches
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_SketchyCOCO_Image_Generation_From_Freehand_Scene_Sketches_CVPR_2020_paper.html
AUTHORS: Chengying Gao, Qi Liu, Qi Xu, Limin Wang, Jianzhuang Liu, Changqing Zou
HIGHLIGHT: We introduce the first method for automatic image generation from scene-level freehand sketches. We have built a large-scale composite dataset called SketchyCOCO to support and evaluate the solution.

513, TITLE: Image Based Virtual Try-On Network From Unpaired Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Neuberger_Image_Based_Virtual_Try-On_Network_From_Unpaired_Data_CVPR_2020_paper.html
AUTHORS: Assaf Neuberger, Eran Borenstein, Bar Hilleli, Eduard Oks, Sharon Alpert
HIGHLIGHT: This paper presents a new image-based virtual try-on approach (Outfit-VITON) that helps visualize how a composition of clothing items selected from various reference images form a cohesive outfit on a person in a query image.

514, TITLE: PSGAN: Pose and Expression Robust Spatial-Aware GAN for Customizable Makeup Transfer
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_PSGAN_Pose_and_Expression_Robust_Spatial-Aware_GAN_for_Customizable_Makeup_CVPR_2020_paper.html
AUTHORS: Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, Shuicheng Yan
HIGHLIGHT: In this paper, we address the makeup transfer task, which aims to transfer the makeup from a reference image to a source image.

515, TITLE: RetinaFace: Single-Shot Multi-Level Face Localisation in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Deng_RetinaFace_Single-Shot_Multi-Level_Face_Localisation_in_the_Wild_CVPR_2020_paper.html
AUTHORS: Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we present a novel single-shot, multi-level face localisation method, named RetinaFace, which unifies face box prediction, 2D facial landmark localisation and 3D vertices regression under one common target: point regression on the image plane.

516, TITLE: Semantic Image Manipulation Using Scene Graphs
http://openaccess.thecvf.com/content_CVPR_2020/html/Dhamo_Semantic_Image_Manipulation_Using_Scene_Graphs_CVPR_2020_paper.html
AUTHORS: Helisa Dhamo, Azade Farshad, Iro Laina, Nassir Navab, Gregory D. Hager, Federico Tombari, Christian Rupprecht
HIGHLIGHT: Our goal is to encode image information in a given constellation and from there on generate new constellations, such as replacing objects or even changing relationships between objects, while respecting the semantics and style from the original image.

517, TITLE: A Stochastic Conditioning Scheme for Diverse Human Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Aliakbarian_A_Stochastic_Conditioning_Scheme_for_Diverse_Human_Motion_Prediction_CVPR_2020_paper.html
AUTHORS: Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, Stephen Gould

HIGHLIGHT: Alternatively, in this paper, we propose to stochastically combine the root of variations with previous pose information, so as to force the model to take the noise into account.

518, TITLE: Transferring Dense Pose to Proximal Animal Classes
http://openaccess.thecvf.com/content_CVPR_2020/html/Sanakoyeu_Transferring_Dense_Pose_to_Proximal_Animal_Classes_CVPR_2020_paper.html
AUTHORS: Artsiom Sanakoyeu, Vasil Khalidov, Maureen S. McCarthy, Andrea Vedaldi, Natalia Neverova
HIGHLIGHT: We show that, at least for proximal animal classes such as chimpanzees, it is possible to transfer the knowledge existing in dense pose recognition for humans, as well as in more general object detectors and segmenters, to the problem of dense pose recognition in other classes.

519, TITLE: Weakly-Supervised 3D Human Pose Learning via Multi-View Images in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Iqbal_Weakly-Supervised_3D_Human_Pose_Learning_via_Multi-View_Images_in_the_CVPR_2020_paper.html
AUTHORS: Umar Iqbal, Pavlo Molchanov, Jan Kautz
HIGHLIGHT: We propose a novel end-to-end learning framework that enables weakly-supervised training using multi-view consistency.

520, TITLE: VIBE: Video Inference for Human Body Pose and Shape Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kocabas_VIBE_Video_Inference_for_Human_Body_Pose_and_Shape_Estimation_CVPR_2020_paper.html
AUTHORS: Muhammed Kocabas, Nikos Athanasiou, Michael J. Black
HIGHLIGHT: To address this problem, we propose "Video Inference for Body Pose and Shape Estimation" (VIBE), which makes use of an existing large-scale motion capture dataset (AMASS) together with unpaired, in-the-wild, 2D keypoint annotations.

521, TITLE: G3AN: Disentangling Appearance and Motion for Video Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_G3AN_Disentangling_Appearance_and_Motion_for_Video_Generation_CVPR_2020_paper.html
AUTHORS: Yaohui Wang, Piotr Bilinski, Francois Bremond, Antitza Dantcheva
HIGHLIGHT: To tackle this challenge, we introduce G3AN, a novel spatio-temporal generative model, which seeks to capture the distribution of high dimensional video data and to model appearance and motion in disentangled manner.

522, TITLE: Domain Adaptive Image-to-Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Domain_Adaptive_Image-to-Image_Translation_CVPR_2020_paper.html
AUTHORS: Ying-Cong Chen, Xiaogang Xu, Jiaya Jia
HIGHLIGHT: To deal with these issues, we propose the Domain Adaptive Image-To-Image translation (DAI2I) framework that adapts an I2I model for out-of-domain samples.

523, TITLE: GAN Compression: Efficient Architectures for Interactive Conditional GANs
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_GAN_Compression_Efficient_Architectures_for_Interactive_Conditional_GANs_CVPR_2020_paper.html
AUTHORS: Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, Song Han
HIGHLIGHT: In this work, we propose a general-purpose compression framework for reducing the inference time and model size of the generator in cGANs.

524, TITLE: Searching Central Difference Convolutional Networks for Face Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Searching_Central_Difference_Convolutional_Networks_for_Face_Anti-Spoofing_CVPR_2020_paper.html
AUTHORS: Zitong Yu, Chenxu Zhao, Zezheng Wang, Yunxiao Qin, Zhuo Su, Xiaobai Li, Feng Zhou, Guoying Zhao
HIGHLIGHT: Here we propose a novel frame level FAS method based on Central Difference Convolution (CDC), which is able to capture intrinsic detailed patterns via aggregating both intensity and gradient information.

525, TITLE: TransMoMo: Invariance-Driven Unsupervised Video Motion Retargeting
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_TransMoMo_Invariance-Driven_Unsupervised_Video_Motion_Retargeting_CVPR_2020_paper.html
AUTHORS: Zhuoqian Yang, Wentao Zhu, Wayne Wu, Chen Qian, Qiang Zhou, Bolei Zhou, Chen Change Loy
HIGHLIGHT: We present a lightweight video motion retargeting approach TransMoMo that is capable of transferring motion of a person in a source video realistically to another video of a target person.

526, TITLE:　　　　　AdaCoF: Adaptive Collaboration of Flows for Video Frame Interpolation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_AdaCoF_Adaptive_Collaboration_of_Flows_for_Video_Frame_Interpolation_CVPR_2020_paper.html
AUTHORS:　　　　　Hyeongmin Lee, Taeoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, Sangyoun Lee
HIGHLIGHT:　　　　　To solve this problem, we propose a new warping module named Adaptive Collaboration of Flows (AdaCoF).

527, TITLE:　　　　　FReeNet: Multi-Identity Face Reenactment
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_FReeNet_Multi-Identity_Face_Reenactment_CVPR_2020_paper.html
AUTHORS:　　　　　Jiangning Zhang, Xianfang Zeng, Mengmeng Wang, Yusu Pan, Liang Liu, Yong Liu, Yu Ding, Changjie Fan
HIGHLIGHT:　　　　　This paper presents a novel multi-identity face reenactment framework, named FReeNet, to transfer facial expressions from an arbitrary source face to a target face with a shared model.

528, TITLE:　　　　　Novel View Synthesis of Dynamic Scenes With Globally Coherent Depths From a Monocular Camera
http://openaccess.thecvf.com/content_CVPR_2020/html/Yoon_Novel_View_Synthesis_of_Dynamic_Scenes_With_Globally_Coherent_Depths_CVPR_2020_paper.html
AUTHORS:　　　　　Jae Shin Yoon, Kihwan Kim, Orazio Gallo, Hyun Soo Park, Jan Kautz
HIGHLIGHT:　　　　　This paper presents a new method to synthesize an image from arbitrary views and times given a collection of images of a dynamic scene.

529, TITLE:　　　　　Monocular Real-Time Hand Shape and Motion Capture Using Multi-Modal Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Monocular_Real-Time_Hand_Shape_and_Motion_Capture_Using_Multi-Modal_Data_CVPR_2020_paper.html
AUTHORS:　　　　　Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, Feng Xu
HIGHLIGHT:　　　　　We present a novel method for monocular hand shape and pose estimation at unprecedented runtime performance of 100fps and at state-of-the-art accuracy.

530, TITLE:　　　　　The GAN That Warped: Semantic Attribute Editing With Unpaired Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Dorta_The_GAN_That_Warped_Semantic_Attribute_Editing_With_Unpaired_Data_CVPR_2020_paper.html
AUTHORS:　　　　　Garoe Dorta, Sara Vicente, Neill D. F. Campbell, Ivor J. A. Simpson
HIGHLIGHT:　　　　　This work proposes to learn how to perform semantic image edits through the application of smooth warp fields.

531, TITLE:　　　　　4D Visualization of Dynamic Events From Unconstrained Multi-View Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Bansal_4D_Visualization_of_Dynamic_Events_From_Unconstrained_Multi-View_Videos_CVPR_2020_paper.html
AUTHORS:　　　　　Aayush Bansal, Minh Vo, Yaser Sheikh, Deva Ramanan, Srinivasa Narasimhan
HIGHLIGHT:　　　　　We present a data-driven approach for 4D space-time visualization of dynamic events from videos captured by hand-held multiple cameras.

532, TITLE:　　　　　Global-Local Bidirectional Reasoning for Unsupervised Representation Learning of 3D Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Rao_Global-Local_Bidirectional_Reasoning_for_Unsupervised_Representation_Learning_of_3D_Point_CVPR_2020_paper.html
AUTHORS:　　　　　Yongming Rao, Jiwen Lu, Jie Zhou
HIGHLIGHT:　　　　　Based on this hypothesis, we propose to learn point cloud representation by bidirectional reasoning between the local structures at different abstraction hierarchies and the global shape without human supervision.

533, TITLE:　　　　　HigherHRNet: Scale-Aware Representation Learning for Bottom-Up Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_HigherHRNet_Scale-Aware_Representation_Learning_for_Bottom-Up_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS:　　　　　Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S. Huang, Lei Zhang
HIGHLIGHT:　　　　　In this paper, we present HigherHRNet: a novel bottom-up human pose estimation method for learning scale-aware representations using high-resolution feature pyramids.

534, TITLE:　　　　　Detecting Attended Visual Targets in Video
http://openaccess.thecvf.com/content_CVPR_2020/html/Chong_Detecting_Attended_Visual_Targets_in_Video_CVPR_2020_paper.html
AUTHORS:　　　　　Eunji Chong, Yongxin Wang, Nataniel Ruiz, James M. Rehg
HIGHLIGHT:　　　　　Our goal is to identify where each person in each frame of a video is looking, and correctly handle the case where the gaze target is out-of-frame.
We introduce a new annotated dataset, VideoAttentionTarget, containing complex and dynamic patterns of real-world gaze behavior.

535, TITLE: Closed-Loop Matters: Dual Regression Networks for Single Image Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Closed-Loop_Matters_Dual_Regression_Networks_for_Single_Image_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Yong Guo, Jian Chen, Jingdong Wang, Qi Chen, Jiezhang Cao, Zeshuai Deng, Yanwu Xu, Mingkui Tan
HIGHLIGHT: To address the above issues, we propose a dual regression scheme by introducing an additional constraint on LR data to reduce the space of the possible functions.

536, TITLE: Neural Voxel Renderer: Learning an Accurate and Controllable Rendering Tool
http://openaccess.thecvf.com/content_CVPR_2020/html/Rematas_Neural_Voxel_Renderer_Learning_an_Accurate_and_Controllable_Rendering_Tool_CVPR_2020_paper.html
AUTHORS: Konstantinos Rematas, Vittorio Ferrari
HIGHLIGHT: We present a neural rendering framework that maps a voxelized scene into a high quality image.

537, TITLE: Neural Contours: Learning to Draw Lines From 3D Shapes
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Neural_Contours_Learning_to_Draw_Lines_From_3D_Shapes_CVPR_2020_paper.html
AUTHORS: Difan Liu, Mohamed Nabail, Aaron Hertzmann, Evangelos Kalogerakis
HIGHLIGHT: This paper introduces a method for learning to generate line drawings from 3D models.

538, TITLE: Softmax Splatting for Video Frame Interpolation
http://openaccess.thecvf.com/content_CVPR_2020/html/Niklaus_Softmax_Splatting_for_Video_Frame_Interpolation_CVPR_2020_paper.html
AUTHORS: Simon Niklaus, Feng Liu
HIGHLIGHT: We propose softmax splatting to address this paradigm shift and show its effectiveness on the application of frame interpolation.

539, TITLE: CIAGAN: Conditional Identity Anonymization Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Maximov_CIAGAN_Conditional_Identity_Anonymization_Generative_Adversarial_Networks_CVPR_2020_paper.html
AUTHORS: Maxim Maximov, Ismail Elezi, Laura Leal-Taixe
HIGHLIGHT: We propose and develop CIAGAN, a model for image and video anonymization based on conditional generative adversarial networks.

540, TITLE: Probabilistic Structural Latent Representation for Unsupervised Embedding
http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_Probabilistic_Structural_Latent_Representation_for_Unsupervised_Embedding_CVPR_2020_paper.html
AUTHORS: Mang Ye, Jianbing Shen
HIGHLIGHT: To tackle these issues, this paper proposes a probabilistic structural latent representation (PSLR), which incorporates an adaptable softmax embedding to approximate the positive concentrated and negative instance separated properties in the graph latent space.

541, TITLE: Semantically Multi-Modal Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Semantically_Multi-Modal_Image_Synthesis_CVPR_2020_paper.html
AUTHORS: Zhen Zhu, Zhiliang Xu, Ansheng You, Xiang Bai
HIGHLIGHT: In this paper, we focus on semantically multi-modal image synthesis (SMIS) task, namely, generating multi-modal images at the semantic level.

542, TITLE: Nested Scale-Editing for Conditional Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Nested_Scale-Editing_for_Conditional_Image_Synthesis_CVPR_2020_paper.html
AUTHORS: Lingzhi Zhang, Jiancong Wang, Yinshuang Xu, Jie Min, Tarmily Wen, James C. Gee, Jianbo Shi
HIGHLIGHT: We propose an image synthesis approach that provides stratified navigation in the latent code space.

543, TITLE: UnrealText: Synthesizing Realistic Scene Text Images From the Unreal World
http://openaccess.thecvf.com/content_CVPR_2020/html/Long_UnrealText_Synthesizing_Realistic_Scene_Text_Images_From_the_Unreal_World_CVPR_2020_paper.html
AUTHORS: Shangbang Long, Cong Yao
HIGHLIGHT: In this paper, we introduce UnrealText, an efficient image synthesis method that renders realistic images via a 3D graphics engine.

544, TITLE: Fast Texture Synthesis via Pseudo Optimizer
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Fast_Texture_Synthesis_via_Pseudo_Optimizer_CVPR_2020_paper.html
AUTHORS: Wu Shi, Yu Qiao
HIGHLIGHT: We propose a new efficient method that aims to simulate the optimization process while retains most of the properties.

545, TITLE: Towards Learning Structure via Consensus for Face Segmentation and Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Masi_Towards_Learning_Structure_via_Consensus_for_Face_Segmentation_and_Parsing_CVPR_2020_paper.html
AUTHORS: Iacopo Masi, Joe Mathai, Wael AbdAlmageed
HIGHLIGHT: We thereby offer a novel learning mechanism to enforce structure in the prediction via consensus, guided by a robust loss function that forces pixel objects to be consistent with each other.

546, TITLE: CookGAN: Causality Based Text-to-Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_CookGAN_Causality_Based_Text-to-Image_Synthesis_CVPR_2020_paper.html
AUTHORS: Bin Zhu, Chong-Wah Ngo
HIGHLIGHT: This paper presents a new network architecture, CookGAN, that mimics visual effect in causality chain, preserves fine-grained details and progressively upsamples image.

547, TITLE: Weakly Supervised Discriminative Feature Learning With State Information for Person Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Weakly_Supervised_Discriminative_Feature_Learning_With_State_Information_for_Person_CVPR_2020_paper.html
AUTHORS: Hong-Xing Yu, Wei-Shi Zheng
HIGHLIGHT: In this work we propose utilizing the state information as weak supervision to address the visual discrepancy caused by different states.

548, TITLE: Future Video Synthesis With Object Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Future_Video_Synthesis_With_Object_Motion_Prediction_CVPR_2020_paper.html
AUTHORS: Yue Wu, Rongrong Gao, Jaesik Park, Qifeng Chen
HIGHLIGHT: We present an approach to predict future video frames given a sequence of continuous video frames in the past.

549, TITLE: MaskGAN: Towards Diverse and Interactive Facial Image Manipulation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_MaskGAN_Towards_Diverse_and_Interactive_Facial_Image_Manipulation_CVPR_2020_paper.html
AUTHORS: Cheng-Han Lee, Ziwei Liu, Lingyun Wu, Ping Luo
HIGHLIGHT: To overcome these drawbacks, we propose a novel framework termed MaskGAN, enabling diverse and interactive face manipulation.
To facilitate extensive studies, we construct a large-scale high-resolution face dataset with fine-grained mask annotations named CelebAMask-HQ.

550, TITLE: A Graduated Filter Method for Large Scale Robust Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Le_A_Graduated_Filter_Method_for_Large_Scale_Robust_Estimation_CVPR_2020_paper.html
AUTHORS: Huu Le, Christopher Zach
HIGHLIGHT: In this paper, we introduce a novel solver for robust estimation that possesses a strong ability to escape poor local minima.

551, TITLE: Deep Face Super-Resolution With Iterative Collaboration Between Attentive Recovery and Landmark Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Ma_Deep_Face_Super-Resolution_With_Iterative_Collaboration_Between_Attentive_Recovery_and_CVPR_2020_paper.html
AUTHORS: Cheng Ma, Zhenyu Jiang, Yongming Rao, Jiwen Lu, Jie Zhou
HIGHLIGHT: In this paper, we propose a deep face super-resolution (FSR) method with iterative collaboration between two recurrent networks which focus on facial image recovery and landmark estimation respectively.

552, TITLE: Coherent Reconstruction of Multiple Humans From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Coherent_Reconstruction_of_Multiple_Humans_From_a_Single_Image_CVPR_2020_paper.html

AUTHORS: Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, Kostas Daniilidis
HIGHLIGHT: In this work, we address the problem of multi-person 3D pose estimation from a single image.

553, TITLE: PointASNL: Robust Point Clouds Processing Using Nonlocal Neural Networks With Adaptive Sampling
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_PointASNL_Robust_Point_Clouds_Processing_Using_Nonlocal_Neural_Networks_With_CVPR_2020_paper.html
AUTHORS: Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, Shuguang Cui
HIGHLIGHT: In this paper, we present a novel end-to-end network for robust point clouds processing, named PointASNL, which can deal with point clouds with noise effectively.

554, TITLE: A Neural Rendering Framework for Free-Viewpoint Relighting
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_A_Neural_Rendering_Framework_for_Free-Viewpoint_Relighting_CVPR_2020_paper.html
AUTHORS: Zhang Chen, Anpei Chen, Guli Zhang, Chengyuan Wang, Yu Ji, Kiriakos N. Kutulakos, Jingyi Yu
HIGHLIGHT: We present a novel Relightable Neural Renderer (RNR) for simultaneous view synthesis and relighting using multi-view image inputs.

555, TITLE: A Multi-Task Mean Teacher for Semi-Supervised Shadow Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_A_Multi-Task_Mean_Teacher_for_Semi-Supervised_Shadow_Detection_CVPR_2020_paper.html
AUTHORS: Zhihao Chen, Lei Zhu, Liang Wan, Song Wang, Wei Feng, Pheng-Ann Heng
HIGHLIGHT: To boost the shadow detection performance, this paper presents a multi-task mean teacher model for semi-supervised shadow detection by leveraging unlabeled data and exploring the learning of multiple information of shadows simultaneously.

556, TITLE: GroupFace: Learning Latent Groups and Constructing Group-Based Representations for Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_GroupFace_Learning_Latent_Groups_and_Constructing_Group-Based_Representations_for_Face_CVPR_2020_paper.html
AUTHORS: Yonghyun Kim, Wonpyo Park, Myung-Cheol Roh, Jongju Shin
HIGHLIGHT: We propose a novel face-recognition-specialized architecture called GroupFace that utilizes multiple group-aware representations, simultaneously, to improve the quality of the embedding feature.

557, TITLE: Channel Attention Based Iterative Residual Learning for Depth Map Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Song_Channel_Attention_Based_Iterative_Residual_Learning_for_Depth_Map_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Xibin Song, Yuchao Dai, Dingfu Zhou, Liu Liu, Wei Li, Hongdong Li, Ruigang Yang
HIGHLIGHT: In this paper, we argue that DSR models trained under this setting are restrictive and not effective in dealing with realworld DSR tasks.

558, TITLE: Time Flies: Animating a Still Image With Time-Lapse Video As Reference
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Time_Flies_Animating_a_Still_Image_With_Time-Lapse_Video_As_CVPR_2020_paper.html
AUTHORS: Chia-Chi Cheng, Hung-Yu Chen, Wei-Chen Chiu
HIGHLIGHT: In this paper, we propose a self-supervised end-to-end model to generate the time-lapse video from a single image and a reference video.

559, TITLE: SER-FIQ: Unsupervised Estimation of Face Image Quality Based on Stochastic Embedding Robustness
http://openaccess.thecvf.com/content_CVPR_2020/html/Terhorst_SER-FIQ_Unsupervised_Estimation_of_Face_Image_Quality_Based_on_Stochastic_CVPR_2020_paper.html
AUTHORS: Philipp Terhorst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, Arjan Kuijper
HIGHLIGHT: Avoiding the use of inaccurate quality labels, we proposed a novel concept to measure face quality based on an arbitrary face recognition model.

560, TITLE: Grid-GCN for Fast and Scalable Point Cloud Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Grid-GCN_for_Fast_and_Scalable_Point_Cloud_Learning_CVPR_2020_paper.html
AUTHORS: Qiangeng Xu, Xudong Sun, Cho-Ying Wu, Panqu Wang, Ulrich Neumann
HIGHLIGHT: In this paper, we present a method, named Grid-GCN, for fast and scalable point cloud learning.

561, TITLE: Domain Balancing: Face Recognition on Long-Tailed Domains

http://openaccess.thecvf.com/content_CVPR_2020/html/Cao_Domain_Balancing_Face_Recognition_on_Long-Tailed_Domains_CVPR_2020_paper.html
AUTHORS: Dong Cao, Xiangyu Zhu, Xingyu Huang, Jianzhu Guo, Zhen Lei
HIGHLIGHT: In this paper, we propose a novel Domain Balancing (DB) mechanism to handle this problem.

562, TITLE: AdversarialNAS: Adversarial Neural Architecture Search for GANs
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_AdversarialNAS_Adversarial_Neural_Architecture_Search_for_GANs_CVPR_2020_paper.html
AUTHORS: Chen Gao, Yunpeng Chen, Si Liu, Zhenxiong Tan, Shuicheng Yan
HIGHLIGHT: In this paper, we propose an AdversarialNAS method specially tailored for Generative Adversarial Networks (GANs) to search for a superior generative model on the task of unconditional image generation.

563, TITLE: Image Super-Resolution With Cross-Scale Non-Local Attention and Exhaustive Self-Exemplars Mining
http://openaccess.thecvf.com/content_CVPR_2020/html/Mei_Image_Super-Resolution_With_Cross-Scale_Non-Local_Attention_and_Exhaustive_Self-Exemplars_Mining_CVPR_2020_paper.html
AUTHORS: Yiqun Mei, Yuchen Fan, Yuqian Zhou, Lichao Huang, Thomas S. Huang, Honghui Shi
HIGHLIGHT: In this paper, we propose the first Cross-Scale Non-Local (CS-NL) attention module with integration into a recurrent neural network.

564, TITLE: The Devil Is in the Details: Delving Into Unbiased Data Processing for Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_The_Devil_Is_in_the_Details_Delving_Into_Unbiased_Data_CVPR_2020_paper.html
AUTHORS: Junjie Huang, Zheng Zhu, Feng Guo, Guan Huang
HIGHLIGHT: In this paper, we focus on this problem and find that the devil of top-down pose estimator is in the biased data processing.

565, TITLE: Data Uncertainty Learning in Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Chang_Data_Uncertainty_Learning_in_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Jie Chang, Zhonghao Lan, Changmao Cheng, Yichen Wei
HIGHLIGHT: This work applies data uncertainty learning to face recognition, such that the feature (mean) and uncertainty (variance) are learnt simultaneously, for the first time.

566, TITLE: Regularizing Discriminative Capability of CGANs for Semi-Supervised Generative Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Regularizing_Discriminative_Capability_of_CGANs_for_Semi-Supervised_Generative_Learning_CVPR_2020_paper.html
AUTHORS: Yi Liu, Guangchang Deng, Xiangping Zeng, Si Wu, Zhiwen Yu, Hau-San Wong
HIGHLIGHT: To address this issue, we propose a regularization technique based on Random Regional Replacement (R^3-regularization) to facilitate the generative learning process.

567, TITLE: FM2u-Net: Face Morphological Multi-Branch Network for Makeup-Invariant Face Verification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_FM2u-Net_Face_Morphological_Multi-Branch_Network_for_Makeup-Invariant_Face_Verification_CVPR_2020_paper.html
AUTHORS: Wenxuan Wang, Yanwei Fu, Xuelin Qian, Yu-Gang Jiang, Qi Tian, Xiangyang Xue
HIGHLIGHT: To address these challenges, we propose a unified Face Morphological Multi-branch Network (FMMu-Net) for makeup-invariant face verification, which can simultaneously synthesize many diverse makeup faces through face morphology network (FM-Net) and effectively learn cosmetics-robust face representations using attention-based multi-branch learning network (AttM-Net).

568, TITLE: UCTGAN: Diverse Image Inpainting Based on Unsupervised Cross-Space Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_UCTGAN_Diverse_Image_Inpainting_Based_on_Unsupervised_Cross-Space_Translation_CVPR_2020_paper.html
AUTHORS: Lei Zhao, Qihang Mo, Sihuan Lin, Zhizhong Wang, Zhiwen Zuo, Haibo Chen, Wei Xing, Dongming Lu
HIGHLIGHT: In order to produce multiple and diverse reasonable solutions, we present Unsupervised Cross-space Translation Generative Adversarial Network (called UCTGAN) which mainly consists of three network modules: conditional encoder module, manifold projection module and generation module.

569, TITLE: Decoupled Representation Learning for Skeleton-Based Gesture Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Decoupled_Representation_Learning_for_Skeleton-Based_Gesture_Recognition_CVPR_2020_paper.html
AUTHORS: Jianbo Liu, Yongcheng Liu, Ying Wang, Veronique Prinet, Shiming Xiang, Chunhong Pan

HIGHLIGHT: In this paper, we propose to decouple the gesture into hand posture variations and hand movements, which are then modeled separately.

570, TITLE: An Efficient PointLSTM for Point Clouds Based Gesture Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Min_An_Efficient_PointLSTM_for_Point_Clouds_Based_Gesture_Recognition_CVPR_2020_paper.html
AUTHORS: Yuecong Min, Yanxiao Zhang, Xiujuan Chai, Xilin Chen
HIGHLIGHT: In this paper, we formulate gesture recognition as an irregular sequence recognition problem and aim to capture long-term spatial correlations across point cloud sequences.

571, TITLE: Editing in Style: Uncovering the Local Semantics of GANs
http://openaccess.thecvf.com/content_CVPR_2020/html/Collins_Editing_in_Style_Uncovering_the_Local_Semantics_of_GANs_CVPR_2020_paper.html
AUTHORS: Edo Collins, Raja Bala, Bob Price, Sabine Susstrunk
HIGHLIGHT: Focusing on StyleGAN, we introduce a simple and effective method for making local, semantically-aware edits to a target output image.

572, TITLE: On the Detection of Digital Face Manipulation
http://openaccess.thecvf.com/content_CVPR_2020/html/Dang_On_the_Detection_of_Digital_Face_Manipulation_CVPR_2020_paper.html
AUTHORS: Hao Dang, Feng Liu, Joel Stehouwer, Xiaoming Liu, Anil K. Jain
HIGHLIGHT: Instead of simply using multi-task learning to simultaneously detect manipulated images and predict the manipulated mask (regions), we propose to utilize an attention mechanism to process and improve the feature maps for the classification task.

573, TITLE: Learning Texture Transformer Network for Image Super-Resolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_Texture_Transformer_Network_for_Image_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, Baining Guo
HIGHLIGHT: In this paper, we propose a novel Texture Transformer Network for Image Super-Resolution (TTSR), in which the LR and Ref images are formulated as queries and keys in a transformer, respectively.

574, TITLE: Reference-Based Sketch Image Colorization Using Augmented-Self Reference and Dense Semantic Correspondence
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Reference-Based_Sketch_Image_Colorization_Using_Augmented-Self_Reference_and_Dense_Semantic_CVPR_2020_paper.html
AUTHORS: Junsoo Lee, Eungyeup Kim, Yunsung Lee, Dongjun Kim, Jaehyuk Chang, Jaegul Choo
HIGHLIGHT: To tackle this challenge, we propose to utilize the identical image with geometric distortion as a virtual reference, which makes it possible to secure the ground truth for a colored output image.

575, TITLE: Deblurring Using Analysis-Synthesis Networks Pair
http://openaccess.thecvf.com/content_CVPR_2020/html/Kaufman_Deblurring_Using_Analysis-Synthesis_Networks_Pair_CVPR_2020_paper.html
AUTHORS: Adam Kaufman, Raanan Fattal
HIGHLIGHT: We propose a new architecture which breaks the deblurring network into an analysis network which estimates the blur, and a synthesis network that uses this kernel to deblur the image.

576, TITLE: Exploring Unlabeled Faces for Novel Attribute Discovery
http://openaccess.thecvf.com/content_CVPR_2020/html/Bahng_Exploring_Unlabeled_Faces_for_Novel_Attribute_Discovery_CVPR_2020_paper.html
AUTHORS: Hyojin Bahng, Sunghyo Chung, Seungjoo Yoo, Jaegul Choo
HIGHLIGHT: In this paper, we attempt to alleviate this necessity for labeled data in the facial image translation domain.

577, TITLE: Neural Pose Transfer by Spatially Adaptive Instance Normalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Neural_Pose_Transfer_by_Spatially_Adaptive_Instance_Normalization_CVPR_2020_paper.html
AUTHORS: Jiashun Wang, Chao Wen, Yanwei Fu, Haitao Lin, Tianyun Zou, Xiangyang Xue, Yinda Zhang
HIGHLIGHT: Particularly in this paper, we are interested in transferring the pose of source human mesh to deform the target human mesh, while the source and target meshes may have different identity information.

578, TITLE: Fine-Grained Image-to-Image Transformation Towards Visual Recognition

http://openaccess.thecvf.com/content_CVPR_2020/html/Xiong_Fine-Grained_Image-to-Image_Transformation_Towards_Visual_Recognition_CVPR_2020_paper.html
AUTHORS: Wei Xiong, Yutong He, Yixuan Zhang, Wenhan Luo, Lin Ma, Jiebo Luo
HIGHLIGHT: In this paper, we aim at transforming an image with a fine-grained category to synthesize new images that preserve the identity of the input image, which can thereby benefit the subsequent fine-grained image recognition and few-shot learning tasks.

579, TITLE: Deep Facial Non-Rigid Multi-View Stereo
http://openaccess.thecvf.com/content_CVPR_2020/html/Bai_Deep_Facial_Non-Rigid_Multi-View_Stereo_CVPR_2020_paper.html
AUTHORS: Ziqian Bai, Zhaopeng Cui, Jamal Ahmed Rahim, Xiaoming Liu, Ping Tan
HIGHLIGHT: We present a method for 3D face reconstruction from multi-view images with different expressions.

580, TITLE: Attention-Driven Cropping for Very High Resolution Facial Landmark Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chandran_Attention-Driven_Cropping_for_Very_High_Resolution_Facial_Landmark_Detection_CVPR_2020_paper.html
AUTHORS: Prashanth Chandran, Derek Bradley, Markus Gross, Thabo Beeler
HIGHLIGHT: Building on top of recent progress in attention-based networks, we present a novel, fully convolutional regional architecture that is specially designed for predicting landmarks on very high resolution facial images without downsampling.

581, TITLE: Towards Unsupervised Learning of Generative Models for 3D Controllable Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Liao_Towards_Unsupervised_Learning_of_Generative_Models_for_3D_Controllable_Image_CVPR_2020_paper.html
AUTHORS: Yiyi Liao, Katja Schwarz, Lars Mescheder, Andreas Geiger
HIGHLIGHT: We define the new task of 3D controllable image synthesis and propose an approach for solving it by reasoning both in 3D space and in the 2D image domain.

582, TITLE: End-to-End Pseudo-LiDAR for Image-Based 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Qian_End-to-End_Pseudo-LiDAR_for_Image-Based_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Belongie, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger, Wei-Lun Chao
HIGHLIGHT: In this paper, we introduce a new framework based on differentiable Change of Representation (CoR) modules that allow the entire PL pipeline to be trained end-to-end.

583, TITLE: Towards High-Fidelity 3D Face Reconstruction From In-the-Wild Images Using Graph Convolutional Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Towards_High-Fidelity_3D_Face_Reconstruction_From_In-the-Wild_Images_Using_Graph_CVPR_2020_paper.html
AUTHORS: Jiangke Lin, Yi Yuan, Tianjia Shao, Kun Zhou
HIGHLIGHT: In this paper, we introduce a method to reconstruct 3D facial shapes with high-fidelity textures from single-view images in the wild, without the need to capture a large-scale face texture database.

584, TITLE: CurricularFace: Adaptive Curriculum Learning Loss for Deep Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_CurricularFace_Adaptive_Curriculum_Learning_Loss_for_Deep_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, Feiyue Huang
HIGHLIGHT: In this work, we propose a novel Adaptive Curriculum Learning loss (CurricularFace) that embeds the idea of curriculum learning into the loss function to achieve a novel training strategy for deep face recognition, which mainly addresses easy samples in the early training stage and hard ones in the later stage.

585, TITLE: Rotate-and-Render: Unsupervised Photorealistic Face Rotation From Single-View Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Rotate-and-Render_Unsupervised_Photorealistic_Face_Rotation_From_Single-View_Images_CVPR_2020_paper.html
AUTHORS: Hang Zhou, Jihao Liu, Ziwei Liu, Yu Liu, Xiaogang Wang
HIGHLIGHT: To overcome these challenges, we propose a novel unsupervised framework that can synthesize photo-realistic rotated faces using only single-view image collections in the wild.

586, TITLE: One-Shot Domain Adaptation for Face Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_One-Shot_Domain_Adaptation_for_Face_Generation_CVPR_2020_paper.html
AUTHORS: Chao Yang, Ser-Nam Lim

HIGHLIGHT: In this paper, we propose a framework capable of generating face images that fall into the same distribution as that of a given one-shot example.

587, TITLE: BidNet: Binocular Image Dehazing Without Explicit Disparity Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Pang_BidNet_Binocular_Image_Dehazing_Without_Explicit_Disparity_Estimation_CVPR_2020_paper.html
AUTHORS: Yanwei Pang, Jing Nie, Jin Xie, Jungong Han, Xuelong Li
HIGHLIGHT: On the assumption that dehazed binocular images are superior to the hazy ones for stereo vision tasks such as 3D object detection and according to the fact that image haze is a function of depth, this paper proposes a Binocular image dehazing Network (BidNet) aiming at dehazing both the left and right images of binocular images within the deep learning framework.

588, TITLE: Deep Shutter Unrolling Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Deep_Shutter_Unrolling_Network_CVPR_2020_paper.html
AUTHORS: Peidong Liu, Zhaopeng Cui, Viktor Larsson, Marc Pollefeys
HIGHLIGHT: We present a novel network for rolling shutter effect correction.

589, TITLE: Joint Texture and Geometry Optimization for RGB-D Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Fu_Joint_Texture_and_Geometry_Optimization_for_RGB-D_Reconstruction_CVPR_2020_paper.html
AUTHORS: Yanping Fu, Qingan Yan, Jie Liao, Chunxia Xiao
HIGHLIGHT: In this paper, we propose a novel approach that can jointly optimize the camera poses, texture and geometry of the reconstructed model, and color consistency between the key-frames.

590, TITLE: Deep 3D Capture: Geometry and Reflectance From Sparse Multi-View Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Bi_Deep_3D_Capture_Geometry_and_Reflectance_From_Sparse_Multi-View_Images_CVPR_2020_paper.html
AUTHORS: Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, Ravi Ramamoorthi
HIGHLIGHT: We introduce a novel learning-based method to reconstruct the high-quality geometry and complex, spatially-varying BRDF of an arbitrary object from a sparse set of only six images captured by wide-baseline cameras under collocated point lighting.

591, TITLE: Auto-Tuning Structured Light by Optical Stochastic Gradient Descent
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Auto-Tuning_Structured_Light_by_Optical_Stochastic_Gradient_Descent_CVPR_2020_paper.html
AUTHORS: Wenzheng Chen, Parsa Mirdehghan, Sanja Fidler, Kiriakos N. Kutulakos
HIGHLIGHT: We consider the problem of optimizing the performance of an active imaging system by automatically discovering the illuminations it should use, and the way to decode them.

592, TITLE: MARMVS: Matching Ambiguity Reduced Multiple View Stereo for Efficient Large Scale Scene Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_MARMVS_Matching_Ambiguity_Reduced_Multiple_View_Stereo_for_Efficient_Large_CVPR_2020_paper.html
AUTHORS: Zhenyu Xu, Yiguang Liu, Xuelei Shi, Ying Wang, Yunan Zheng
HIGHLIGHT: In this paper, we present a novel method, matching ambiguity reduced multiple view stereo (MARMVS) to address this issue.

593, TITLE: Uncertainty Based Camera Model Selection
http://openaccess.thecvf.com/content_CVPR_2020/html/Polic_Uncertainty_Based_Camera_Model_Selection_CVPR_2020_paper.html
AUTHORS: Michal Polic, Stanislav Steidl, Cenek Albl, Zuzana Kukelova, Tomas Pajdla
HIGHLIGHT: In this paper, we present a new automatic method for camera model selection in large scale SfM that is based on efficient uncertainty evaluation.

594, TITLE: Local Implicit Grid Representations for 3D Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Local_Implicit_Grid_Representations_for_3D_Scenes_CVPR_2020_paper.html
AUTHORS: Chiyu "Max" Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Niessner, Thomas Funkhouser
HIGHLIGHT: In this paper, we introduce Local Implicit Grid Representations, a new 3D shape representation designed for scalability and generality.

595, TITLE: TetraTSDF: 3D Human Reconstruction From a Single Image With a Tetrahedral Outer Shell
http://openaccess.thecvf.com/content_CVPR_2020/html/Onizuka_TetraTSDF_3D_Human_Reconstruction_From_a_Single_Image_With_a_CVPR_2020_paper.html
AUTHORS: Hayato Onizuka, Zehra Hayirci, Diego Thomas, Akihiro Sugimoto, Hideaki Uchiyama, Rin-ichiro Taniguchi
HIGHLIGHT: In this paper, we propose the tetrahedral outer shell volumetric truncated signed distance function (TetraTSDF) model for the human body, and its corresponding part connection network (PCN) for 3D human body shape regression.

596, TITLE: Averaging Essential and Fundamental Matrices in Collinear Camera Settings
http://openaccess.thecvf.com/content_CVPR_2020/html/Geifman_Averaging_Essential_and_Fundamental_Matrices_in_Collinear_Camera_Settings_CVPR_2020_paper.html
AUTHORS: Amnon Geifman, Yoni Kasten, Meirav Galun, Ronen Basri
HIGHLIGHT: In this paper, we introduce an analysis and algorithms for averaging bifocal tensors (essential or fundamental matrices) when either subsets or all of the camera centers are collinear.

597, TITLE: On the Distribution of Minima in Intrinsic-Metric Rotation Averaging
http://openaccess.thecvf.com/content_CVPR_2020/html/Wilson_On_the_Distribution_of_Minima_in_Intrinsic-Metric_Rotation_Averaging_CVPR_2020_paper.html
AUTHORS: Kyle Wilson, David Bindel
HIGHLIGHT: In this paper, we study the spatial distribution of local minima.

598, TITLE: Lightweight Multi-View 3D Pose Estimation Through Camera-Disentangled Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Remelli_Lightweight_Multi-View_3D_Pose_Estimation_Through_Camera-Disentangled_Representation_CVPR_2020_paper.html
AUTHORS: Edoardo Remelli, Shangchen Han, Sina Honari, Pascal Fua, Robert Wang
HIGHLIGHT: We present a lightweight solution to recover 3D pose from multi-view images captured with spatially calibrated cameras.

599, TITLE: A Novel Recurrent Encoder-Decoder Structure for Large-Scale Multi-View Stereo Reconstruction From an Open Aerial Dataset
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_A_Novel_Recurrent_Encoder-Decoder_Structure_for_Large-Scale_Multi-View_Stereo_Reconstruction_CVPR_2020_paper.html
AUTHORS: Jin Liu, Shunping Ji
HIGHLIGHT: We also introduce in this paper a novel network, called RED-Net, for wide-range depth inference, which we developed from a recurrent encoder-decoder structure to regularize cost maps across depths and a 2D fully convolutional network as framework.

600, TITLE: Factorized Higher-Order CNNs With an Application to Spatio-Temporal Emotion Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kossaifi_Factorized_Higher-Order_CNNs_With_an_Application_to_Spatio-Temporal_Emotion_Estimation_CVPR_2020_paper.html
AUTHORS: Jean Kossaifi, Antoine Toisoul, Adrian Bulat, Yannis Panagakis, Timothy M. Hospedales, Maja Pantic
HIGHLIGHT: In this paper, we unify these two approaches by proposing a tensor factorization framework for efficient multidimensional (separable) convolutions of higher-order.

601, TITLE: Effectively Unbiased FID and Inception Score and Where to Find Them
http://openaccess.thecvf.com/content_CVPR_2020/html/Chong_Effectively_Unbiased_FID_and_Inception_Score_and_Where_to_Find_CVPR_2020_paper.html
AUTHORS: Min Jin Chong, David Forsyth
HIGHLIGHT: This paper shows that two commonly used evaluation metrics for generative models, the Frechet Inception Distance (FID) and the Inception Score (IS), are biased -- the expected value of the score computed for a finite sample set is not the true value of the score.

602, TITLE: Robust Homography Estimation via Dual Principal Component Pursuit
http://openaccess.thecvf.com/content_CVPR_2020/html/Ding_Robust_Homography_Estimation_via_Dual_Principal_Component_Pursuit_CVPR_2020_paper.html
AUTHORS: Tianjiao Ding, Yunchen Yang, Zhihui Zhu, Daniel P. Robinson, Rene Vidal, Laurent Kneip, Manolis C. Tsakiris
HIGHLIGHT: We revisit robust estimation of homographies over point correspondences between two or three views, a fundamental problem in geometric vision.

603, TITLE: Non-Adversarial Video Synthesis With Learned Priors
http://openaccess.thecvf.com/content_CVPR_2020/html/Aich_Non-Adversarial_Video_Synthesis_With_Learned_Priors_CVPR_2020_paper.html

AUTHORS: Abhishek Aich, Akash Gupta, Rameswar Panda, Rakib Hyder, M. Salman Asif, Amit K. Roy-Chowdhury
HIGHLIGHT: Different from these methods, we focus on the problem of generating videos from latent noise vectors, without any reference input frames.

604, TITLE: Uncertainty-Aware Mesh Decoder for High Fidelity 3D Face Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Uncertainty-
Aware_Mesh_Decoder_for_High_Fidelity_3D_Face_Reconstruction_CVPR_2020_paper.html
AUTHORS: Gun-Hee Lee, Seong-Whan Lee
HIGHLIGHT: In this paper, we propose to employ (i) an uncertainty-aware encoder that presents face features as distributions and (ii) a fully nonlinear decoder model combining Graph CNN with GAN.

605, TITLE: 3FabRec: Fast Few-Shot Face Alignment by Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Browatzki_3FabRec_Fast_Few-
Shot_Face_Alignment_by_Reconstruction_CVPR_2020_paper.html
AUTHORS: Bjorn Browatzki, Christian Wallraven
HIGHLIGHT: We introduce a semi-supervised method in which the crucial idea is to first generate implicit face knowledge from the large amounts of unlabeled images of faces available today.

606, TITLE: Weakly-Supervised Domain Adaptation via GAN and Mesh Model for Estimating 3D Hand Poses Interacting Objects
http://openaccess.thecvf.com/content_CVPR_2020/html/Baek_Weakly-
Supervised_Domain_Adaptation_via_GAN_and_Mesh_Model_for_Estimating_CVPR_2020_paper.html
AUTHORS: Seungryul Baek, Kwang In Kim, Tae-Kyun Kim
HIGHLIGHT: In this work, we propose a novel end-to-end trainable pipeline that adapts the hand-object domain to the single hand-only domain, while learning for HPE.

607, TITLE: Vec2Face: Unveil Human Faces From Their Blackbox Features in Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Duong_Vec2Face_Unveil_Human_Faces_From_Their_Blackbox_Features_in_Face_CVPR_2020_paper.html
AUTHORS: Chi Nhan Duong, Thanh-Dat Truong, Khoa Luu, Kha Gia Quach, Hung Bui, Kaushik Roy
HIGHLIGHT: This paper presents a novel generative structure with Bijective Metric Learning, namely Bijective Generative Adversarial Networks in a Distillation framework (DiBiGAN), for synthesizing faces of an identity given that person's features.

608, TITLE: StyleRig: Rigging StyleGAN for 3D Control Over Portrait Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Tewari_StyleRig_Rigging_StyleGAN_for_3D_Control_Over_Portrait_Images_CVPR_2020_paper.html
AUTHORS: Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Perez, Michael Zollhofer, Christian Theobalt
HIGHLIGHT: We present the first method to provide a face rig-like control over a pretrained and fixed StyleGAN via a 3DMM.

609, TITLE: Self-Supervised 3D Human Pose Estimation via Part Guided Novel Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Kundu_Self-
Supervised_3D_Human_Pose_Estimation_via_Part_Guided_Novel_Image_CVPR_2020_paper.html
AUTHORS: Jogendra Nath Kundu, Siddharth Seth, Varun Jampani, Mugalodi Rakesh, R. Venkatesh Babu, Anirban Chakraborty
HIGHLIGHT: Acknowledging this, we propose a self-supervised learning framework to disentangle such variations from unlabeled video frames.

610, TITLE: Learning Meta Face Recognition in Unseen Domains
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Learning_Meta_Face_Recognition_in_Unseen_Domains_CVPR_2020_paper.html
AUTHORS: Jianzhu Guo, Xiangyu Zhu, Chenxu Zhao, Dong Cao, Zhen Lei, Stan Z. Li
HIGHLIGHT: In this paper, we aim to learn a generalized model that can directly handle new unseen domains without any model updating.
Besides, we propose two benchmarks for generalized face recognition evaluation.

611, TITLE: Cascaded Deep Monocular 3D Human Pose Estimation With Evolutionary Training Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Cascaded_Deep_Monocular_3D_Human_Pose_Estimation_With_Evolutionary_Training_CVPR_2020_paper.html
AUTHORS: Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, Kwang-Ting Cheng

70

HIGHLIGHT: This paper proposes a novel data augmentation method that: (1) is scalable for synthesizing massive amount of training data (over 8 million valid 3D human poses with corresponding 2D projections) for training 2D-to-3D networks, (2) can effectively reduce dataset bias.

612, TITLE: GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_GHUM__GHUML_Generative_3D_Human_Shape_and_Articulated_Pose_CVPR_2020_paper.html
AUTHORS: Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T. Freeman, Rahul Sukthankar, Cristian Sminchisescu
HIGHLIGHT: We present a statistical, articulated 3D human shape modeling pipeline, within a fully trainable, modular, deep learning framework.

613, TITLE: Generating 3D People in Scenes Without People
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Generating_3D_People_in_Scenes_Without_People_CVPR_2020_paper.html
AUTHORS: Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J. Black, Siyu Tang
HIGHLIGHT: We present a fully automatic system that takes a 3D scene and generates plausible 3D human bodies that are posed naturally in that 3D scene.

614, TITLE: Transferring Cross-Domain Knowledge for Video Sign Language Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Transferring_Cross-Domain_Knowledge_for_Video_Sign_Language_Recognition_CVPR_2020_paper.html
AUTHORS: Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, Hongdong Li
HIGHLIGHT: Motivated by this observation, we propose a novel method that learns domain-invariant visual concepts and fertilizes WSLR models by transferring knowledge of subtitled news sign to them.

615, TITLE: Bodies at Rest: 3D Human Pose and Shape Estimation From a Pressure Image Using Synthetic Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Clever_Bodies_at_Rest_3D_Human_Pose_and_Shape_Estimation_From_CVPR_2020_paper.html
AUTHORS: Henry M. Clever, Zackory Erickson, Ariel Kapusta, Greg Turk, Karen Liu, Charles C. Kemp
HIGHLIGHT: We describe a physics-based method that simulates human bodies at rest in a bed with a pressure sensing mat, and present PressurePose, a synthetic dataset with 206K pressure images with 3D human poses and shapes.

616, TITLE: Bayesian Adversarial Human Motion Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Bayesian_Adversarial_Human_Motion_Synthesis_CVPR_2020_paper.html
AUTHORS: Rui Zhao, Hui Su, Qiang Ji
HIGHLIGHT: We propose a generative probabilistic model for human motion synthesis.

617, TITLE: LSM: Learning Subspace Minimization for Low-Level Vision
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_LSM_Learning_Subspace_Minimization_for_Low-Level_Vision_CVPR_2020_paper.html
AUTHORS: Chengzhou Tang, Lu Yuan, Ping Tan
HIGHLIGHT: We study the energy minimization problem in low-level vision tasks from a novel perspective.

618, TITLE: Learning a Neural Solver for Multiple Object Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Braso_Learning_a_Neural_Solver_for_Multiple_Object_Tracking_CVPR_2020_paper.html
AUTHORS: Guillem Braso, Laura Leal-Taixe
HIGHLIGHT: In this work, we exploit the classical network flow formulation of MOT to define a fully differentiable framework based on Message Passing Networks (MPNs).

619, TITLE: GLU-Net: Global-Local Universal Network for Dense Flow and Correspondences
http://openaccess.thecvf.com/content_CVPR_2020/html/Truong_GLU-Net_Global-Local_Universal_Network_for_Dense_Flow_and_Correspondences_CVPR_2020_paper.html
AUTHORS: Prune Truong, Martin Danelljan, Radu Timofte
HIGHLIGHT: In this work, we propose a universal network architecture that is directly applicable to all the aforementioned dense correspondence problems.

620, TITLE: SiamCAR: Siamese Fully Convolutional Classification and Regression for Visual Tracking

http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_SiamCAR_Siamese_Fully_Convolutional_Classification_and_Regression_for_Visual_Tracking_CVPR_2020_paper.html
AUTHORS: Dongyan Guo, Jun Wang, Ying Cui, Zhenhua Wang, Shengyong Chen
HIGHLIGHT: By decomposing the visual tracking task into two subproblems as classification for pixel category and regression for object bounding box at this pixel, we propose a novel fully convolutional Siamese network to solve visual tracking end-to-end in a per-pixel manner.

621, TITLE: MaskFlownet: Asymmetric Feature Matching With Learnable Occlusion Mask
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_MaskFlownet_Asymmetric_Feature_Matching_With_Learnable_Occlusion_Mask_CVPR_2020_paper.html
AUTHORS: Shengyu Zhao, Yilun Sheng, Yue Dong, Eric I-Chao Chang, Yan Xu
HIGHLIGHT: In this paper, we propose an asymmetric occlusion-aware feature matching module, which can learn a rough occlusion mask that filters useless (occluded) areas immediately after feature warping without any explicit supervision.

622, TITLE: Tracking by Instance Detection: A Meta-Learning Approach
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Tracking_by_Instance_Detection_A_Meta-Learning_Approach_CVPR_2020_paper.html
AUTHORS: Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, Wenjun Zeng
HIGHLIGHT: We propose a principled three-step approach to build a high-performance tracker.

623, TITLE: High-Performance Long-Term Tracking With Meta-Updater
http://openaccess.thecvf.com/content_CVPR_2020/html/Dai_High-Performance_Long-Term_Tracking_With_Meta-Updater_CVPR_2020_paper.html
AUTHORS: Kenan Dai, Yunhua Zhang, Dong Wang, Jianhua Li, Huchuan Lu, Xiaoyun Yang
HIGHLIGHT: In this work, we propose a novel offline-trained Meta-Updater to address an important but unsolved problem: Is the tracker ready for updating in the current frame?

624, TITLE: TubeTK: Adopting Tubes to Track Multi-Object in a One-Step Training Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Pang_TubeTK_Adopting_Tubes_to_Track_Multi-Object_in_a_One-Step_Training_CVPR_2020_paper.html
AUTHORS: Bo Pang, Yizhuo Li, Yifan Zhang, Muchen Li, Cewu Lu
HIGHLIGHT: To address these challenges, we propose a concise end-to-end model TubeTK which only needs one step training by introducing the "bounding-tube" to indicate temporal-spatial locations of objects in a short video clip.

625, TITLE: Collaborative Motion Prediction via Neural Motion Message Passing
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Collaborative_Motion_Prediction_via_Neural_Motion_Message_Passing_CVPR_2020_paper.html
AUTHORS: Yue Hu, Siheng Chen, Ya Zhang, Xiao Gu
HIGHLIGHT: To address this challenge, we propose neural motion message passing (NMMP) to explicitly model the interaction and learn representations for directed interactions between actors.

626, TITLE: P2B: Point-to-Box Network for 3D Object Tracking in Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Qi_P2B_Point-to-Box_Network_for_3D_Object_Tracking_in_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Haozhe Qi, Chen Feng, Zhiguo Cao, Feng Zhao, Yang Xiao
HIGHLIGHT: Our main idea is to first localize potential target centers in 3D search area embedded with target information. Then point-driven 3D target proposal and verification are executed jointly.

627, TITLE: Self-Supervised Deep Visual Odometry With Online Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Self-Supervised_Deep_Visual_Odometry_With_Online_Adaptation_CVPR_2020_paper.html
AUTHORS: Shunkai Li, Xin Wang, Yingdian Cao, Fei Xue, Zike Yan, Hongbin Zha
HIGHLIGHT: In this paper, we propose an online meta-learning algorithm to enable VO networks to continuously adapt to new environments in a self-supervised manner.

628, TITLE: Globally Optimal Contrast Maximisation for Event-Based Motion Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Globally_Optimal_Contrast_Maximisation_for_Event-Based_Motion_Estimation_CVPR_2020_paper.html
AUTHORS: Daqi Liu, Alvaro Parra, Tat-Jun Chin
HIGHLIGHT: To alleviate this weakness, we propose a new globally optimal event-based motion estimation algorithm.

629, TITLE:	D3Feat: Joint Learning of Dense Detection and Description of 3D Local Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Bai_D3Feat_Joint_Learning_of_Dense_Detection_and_Description_of_3D_
CVPR_2020_paper.html
AUTHORS:	Xuyang Bai, Zixin Luo, Lei Zhou, Hongbo Fu, Long Quan, Chiew-Lan Tai
HIGHLIGHT:	In this paper, we leverage a 3D fully convolutional network for 3D point clouds, and propose a novel and
practical learning mechanism that densely predicts both a detection score and a description feature for each 3D point.

630, TITLE:	Towards Backward-Compatible Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Towards_Backward-
Compatible_Representation_Learning_CVPR_2020_paper.html
AUTHORS:	Yantao Shen, Yuanjun Xiong, Wei Xia, Stefano Soatto
HIGHLIGHT:	We propose a framework to train embedding models, called backward-compatible training (BCT), as a first step
towards backward compatible representation learning.

631, TITLE:	PointAugment: An Auto-Augmentation Framework for Point Cloud Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_PointAugment_An_Auto-
Augmentation_Framework_for_Point_Cloud_Classification_CVPR_2020_paper.html
AUTHORS:	Ruihui Li, Xianzhi Li, Pheng-Ann Heng, Chi-Wing Fu
HIGHLIGHT:	We present PointAugment, a new auto-augmentation framework that automatically optimizes and augments
point cloud samples to enrich the data diversity when we train a classification network.

632, TITLE:	Cross-Batch Memory for Embedding Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Cross-
Batch_Memory_for_Embedding_Learning_CVPR_2020_paper.html
AUTHORS:	Xun Wang, Haozhi Zhang, Weilin Huang, Matthew R. Scott
HIGHLIGHT:	In this paper, we identify a "slow drift" phenomena by observing that the embedding features drift
exceptionally slow even as the model parameters are updating throughout the training process.

633, TITLE:	Circle Loss: A Unified Perspective of Pair Similarity Optimization
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Circle_Loss_A_Unified_Perspective_of_Pair_Similarity_Optimization_
CVPR_2020_paper.html
AUTHORS:	Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, Yichen Wei
HIGHLIGHT:	This paper provides a pair similarity optimization viewpoint on deep feature learning, aiming to maximize the
within-class similarity s_p and minimize the between-class similarity s_n.

634, TITLE:	Steering Self-Supervised Feature Learning Beyond Local Pixel Statistics
http://openaccess.thecvf.com/content_CVPR_2020/html/Jenni_Steering_Self-
Supervised_Feature_Learning_Beyond_Local_Pixel_Statistics_CVPR_2020_paper.html
AUTHORS:	Simon Jenni, Hailin Jin, Paolo Favaro
HIGHLIGHT:	We introduce a novel principle for self-supervised feature learning based on the discrimination of specific
transformations of an image.

635, TITLE:	Hyperbolic Image Embeddings
http://openaccess.thecvf.com/content_CVPR_2020/html/Khrulkov_Hyperbolic_Image_Embeddings_CVPR_2020_paper.html
AUTHORS:	Valentin Khrulkov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, Victor Lempitsky
HIGHLIGHT:	In this work, we demonstrate that in many practical scenarios, hyperbolic embeddings provide a better
alternative.

636, TITLE:	Controllable Orthogonalization in Training DNNs
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Controllable_Orthogonalization_in_Training_DNNs_CVPR_2020_p
aper.html
AUTHORS:	Lei Huang, Li Liu, Fan Zhu, Diwen Wan, Zehuan Yuan, Bo Li, Ling Shao
HIGHLIGHT:	This paper proposes a computationally efficient and numerically stable orthogonalization method using
Newton's iteration (ONI), to learn a layer-wise orthogonal weight matrix in DNNs.

637, TITLE:	An Investigation Into the Stochasticity of Batch Whitening
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_An_Investigation_Into_the_Stochasticity_of_Batch_Whitening_CVP
R_2020_paper.html
AUTHORS:	Lei Huang, Lei Zhao, Yi Zhou, Fan Zhu, Li Liu, Ling Shao
HIGHLIGHT:	Based on our analysis, we provide a framework for designing and comparing BW algorithms in different
scenarios.

638, TITLE:	High-Order Information Matters: Learning Relation and Topology for Occluded Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_High-
Order_Information_Matters_Learning_Relation_and_Topology_for_Occluded_Person_CVPR_2020_paper.html
AUTHORS:	Guan'an Wang, Shuo Yang, Huanyu Liu, Zhicheng Wang, Yang Yang, Shuliang Wang, Gang Yu, Erjin Zhou, Jian Sun
HIGHLIGHT:	In this paper, we propose a novel framework by learning high-order relation and topology information for discriminative features and robust alignment.

639, TITLE:	Same Features, Different Day: Weakly Supervised Feature Learning for Seasonal Invariance
http://openaccess.thecvf.com/content_CVPR_2020/html/Spencer_Same_Features_Different_Day_Weakly_Supervised_Feature_Learning_for_Seasonal_CVPR_2020_paper.html
AUTHORS:	Jaime Spencer, Richard Bowden, Simon Hadfield
HIGHLIGHT:	The aim of this paper is to provide a dense feature representation that can be used to perform localization, sparse matching or image retrieval, regardless of the current seasonal or temporal appearance.

640, TITLE:	Learning to Dress 3D People in Generative Clothing
http://openaccess.thecvf.com/content_CVPR_2020/html/Ma_Learning_to_Dress_3D_People_in_Generative_Clothing_CVPR_2020_paper.html
AUTHORS:	Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, Michael J. Black
HIGHLIGHT:	To address this, we learn a generative 3D mesh model of clothed people from 3D scans with varying pose and clothing.

641, TITLE:	MAST: A Memory-Augmented Self-Supervised Tracker
http://openaccess.thecvf.com/content_CVPR_2020/html/Lai_MAST_A_Memory-Augmented_Self-Supervised_Tracker_CVPR_2020_paper.html
AUTHORS:	Zihang Lai, Erika Lu, Weidi Xie
HIGHLIGHT:	We propose a dense tracking model trained on videos without any annotations that surpasses previous self-supervised methods on existing benchmarks by a significant margin (+15%), and achieves performance comparable to supervised methods.

642, TITLE:	Learning by Analogy: Reliable Supervision From Transformations for Unsupervised Optical Flow Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Learning_by_Analogy_Reliable_Supervision_From_Transformations_for_Unsupervised_Optical_CVPR_2020_paper.html
AUTHORS:	Liang Liu, Jiangning Zhang, Ruifei He, Yong Liu, Yabiao Wang, Ying Tai, Donghao Luo, Chengjie Wang, Jilin Li, Feiyue Huang
HIGHLIGHT:	In this work, we present a framework to use more reliable supervision from transformations.

643, TITLE:	GNN3DMOT: Graph Neural Network for 3D Multi-Object Tracking With 2D-3D Multi-Feature Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Weng_GNN3DMOT_Graph_Neural_Network_for_3D_Multi-Object_Tracking_With_2D-3D_CVPR_2020_paper.html
AUTHORS:	Xinshuo Weng, Yongxin Wang, Yunze Man, Kris M. Kitani
HIGHLIGHT:	In this work, we propose two techniques to improve the discriminative feature learning for MOT: (1) instead of obtaining features for each object independently, we propose a novel feature interaction mechanism by introducing the Graph Neural Network.

644, TITLE:	ClusterFit: Improving Generalization of Visual Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_ClusterFit_Improving_Generalization_of_Visual_Representations_CVPR_2020_paper.html
AUTHORS:	Xueting Yan, Ishan Misra, Abhinav Gupta, Deepti Ghadiyaram, Dhruv Mahajan
HIGHLIGHT:	In this work, we present a simple strategy - ClusterFit to improve the robustness of the visual representations learned during pre-training.

645, TITLE:	Learning Dynamic Relationships for 3D Human Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Cui_Learning_Dynamic_Relationships_for_3D_Human_Motion_Prediction_CVPR_2020_paper.html
AUTHORS:	Qiongjie Cui, Huaijiang Sun, Fei Yang
HIGHLIGHT:	To tackle these issues, we propose a deep generative model based on graph networks and adversarial learning.

646, TITLE:	Knowledge As Priors: Cross-Modal Knowledge Generalization for Datasets Without Superior Knowledge
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Knowledge_As_Priors_Cross-Modal_Knowledge_Generalization_for_Datasets_Without_Superior_CVPR_2020_paper.html

AUTHORS: Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, Dimitris N. Metaxas
HIGHLIGHT: In this paper, we propose a novel scheme to train the Student in a Target dataset where the Teacher is unavailable.

647, TITLE: S3VAE: Self-Supervised Sequential VAE for Representation Disentanglement and Data Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_S3VAE_Self-Supervised_Sequential_VAE_for_Representation_Disentanglement_and_Data_Generation_CVPR_2020_paper.html
AUTHORS: Yizhe Zhu, Martin Renqiang Min, Asim Kadav, Hans Peter Graf
HIGHLIGHT: We propose a sequential variational autoencoder to learn disentangled representations of sequential data (e.g., videos and audios) under self-supervision.

648, TITLE: Video Playback Rate Perception for Self-Supervised Spatio-Temporal Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yao_Video_Playback_Rate_Perception_for_Self-Supervised_Spatio-Temporal_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Yuan Yao, Chang Liu, Dezhao Luo, Yu Zhou, Qixiang Ye
HIGHLIGHT: In this paper, we propose a novel self-supervised method, referred to as video Playback Rate Perception (PRP), to learn spatio-temporal representation in a simple-yet-effective way.

649, TITLE: Learning to Manipulate Individual Objects in an Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_to_Manipulate_Individual_Objects_in_an_Image_CVPR_2020_paper.html
AUTHORS: Yanchao Yang, Yutong Chen, Stefano Soatto
HIGHLIGHT: The key to our method is the combination of spatial disentanglement, enforced by a Contextual Information Separation loss, and perceptual cycle-consistency, enforced by a loss that penalizes changes in the image partition in response to perturbations of the latent factors.

650, TITLE: PADS: Policy-Adapted Sampling for Visual Similarity Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Roth_PADS_Policy-Adapted_Sampling_for_Visual_Similarity_Learning_CVPR_2020_paper.html
AUTHORS: Karsten Roth, Timo Milbich, Bjorn Ommer
HIGHLIGHT: We, therefore, employ reinforcement learning and have a teacher network adjust the sampling distribution based on the current state of the learner network, which represents visual similarity.

651, TITLE: Siam R-CNN: Visual Tracking by Re-Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Voigtlaender_Siam_R-CNN_Visual_Tracking_by_Re-Detection_CVPR_2020_paper.html
AUTHORS: Paul Voigtlaender, Jonathon Luiten, Philip H.S. Torr, Bastian Leibe
HIGHLIGHT: We present Siam R-CNN, a Siamese re-detection architecture which unleashes the full power of two-stage object detection approaches for visual object tracking.

652, TITLE: ASLFeat: Learning Local Features of Accurate Shape and Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_ASLFeat_Learning_Local_Features_of_Accurate_Shape_and_Localization_CVPR_2020_paper.html
AUTHORS: Zixin Luo, Lei Zhou, Xuyang Bai, Hongkai Chen, Jiahui Zhang, Yao Yao, Shiwei Li, Tian Fang, Long Quan
HIGHLIGHT: In this paper, we present ASLFeat, with three light-weight yet effective modifications to mitigate above issues.

653, TITLE: Filter Grafting for Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Meng_Filter_Grafting_for_Deep_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Fanxu Meng, Hao Cheng, Ke Li, Zhixin Xu, Rongrong Ji, Xing Sun, Guangming Lu
HIGHLIGHT: This paper proposes a new learning paradigm called filter grafting, which aims to improve the representation capability of Deep Neural Networks (DNNs).

654, TITLE: HOPE-Net: A Graph-Based Model for Hand-Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Doosti_HOPE-Net_A_Graph-Based_Model_for_Hand-Object_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Bardia Doosti, Shujon Naha, Majid Mirbagheri, David J. Crandall
HIGHLIGHT: In this paper, we propose a lightweight model called HOPE-Net which jointly estimates hand and object pose in 2D and 3D in real-time.

655, TITLE: DeepFaceFlow: In-the-Wild Dense 3D Facial Motion Estimation

http://openaccess.thecvf.com/content_CVPR_2020/html/Koujan_DeepFaceFlow_In-the-Wild_Dense_3D_Facial_Motion_Estimation_CVPR_2020_paper.html
AUTHORS: Mohammad Rami Koujan, Anastasios Roussos, Stefanos Zafeiriou
HIGHLIGHT: In this work, we propose DeepFaceFlow, a robust, fast, and highly-accurate framework for the dense estimation of 3D non-rigid facial flow between pairs of monocular images.

656, TITLE: Learning for Video Compression With Hierarchical Quality and Recurrent Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_for_Video_Compression_With_Hierarchical_Quality_and_Recurrent_Enhancement_CVPR_2020_paper.html
AUTHORS: Ren Yang, Fabian Mentzer, Luc Van Gool, Radu Timofte
HIGHLIGHT: In this paper, we propose a Hierarchical Learned Video Compression (HLVC) method with three hierarchical quality layers and a recurrent enhancement network.

657, TITLE: Learning Better Lossless Compression Using Lossy Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Mentzer_Learning_Better_Lossless_Compression_Using_Lossy_Compression_CVPR_2020_paper.html
AUTHORS: Fabian Mentzer, Luc Van Gool, Michael Tschannen
HIGHLIGHT: We leverage the powerful lossy image compression algorithm BPG to build a lossless image compression system.

658, TITLE: Flow2Stereo: Effective Self-Supervised Learning of Optical Flow and Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Flow2Stereo_Effective_Self-Supervised_Learning_of_Optical_Flow_and_Stereo_Matching_CVPR_2020_paper.html
AUTHORS: Pengpeng Liu, Irwin King, Michael R. Lyu, Jia Xu
HIGHLIGHT: In this paper, we propose a unified method to jointly learn optical flow and stereo matching.

659, TITLE: Multi-Scale Fusion Subspace Clustering Using Similarity Constraint
http://openaccess.thecvf.com/content_CVPR_2020/html/Dang_Multi-Scale_Fusion_Subspace_Clustering_Using_Similarity_Constraint_CVPR_2020_paper.html
AUTHORS: Zhiyuan Dang, Cheng Deng, Xu Yang, Heng Huang
HIGHLIGHT: In this paper, we propose the Multi-Scale Fusion Subspace Clustering Using Similarity Constraint (SC-MSFSC) network, which learns a more discriminative self-expression coefficient matrix by a novel multi-scale fusion module.

660, TITLE: Siamese Box Adaptive Network for Visual Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Siamese_Box_Adaptive_Network_for_Visual_Tracking_CVPR_2020_paper.html
AUTHORS: Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, Rongrong Ji
HIGHLIGHT: To address this issue, we propose simple yet effective visual tracking framework (named Siamese Box Adaptive Network, SiamBAN) by exploiting the expressive power of the fully convolutional network (FCN).

661, TITLE: Cross-Domain Face Presentation Attack Detection via Multi-Domain Disentangled Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Cross-Domain_Face_Presentation_Attack_Detection_via_Multi-Domain_Disentangled_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Guoqing Wang, Hu Han, Shiguang Shan, Xilin Chen
HIGHLIGHT: In light of this, we propose an efficient disentangled representation learning for cross-domain face PAD.

662, TITLE: Online Deep Clustering for Unsupervised Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhan_Online_Deep_Clustering_for_Unsupervised_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Xiaohang Zhan, Jiahao Xie, Ziwei Liu, Yew-Soon Ong, Chen Change Loy
HIGHLIGHT: To overcome this challenge, we propose Online Deep Clustering (ODC) that performs clustering and network update simultaneously rather than alternatingly.

663, TITLE: Density-Aware Feature Embedding for Face Clustering
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Density-Aware_Feature_Embedding_for_Face_Clustering_CVPR_2020_paper.html
AUTHORS: Senhui Guo, Jing Xu, Dapeng Chen, Chao Zhang, Xiaogang Wang, Rui Zhao
HIGHLIGHT: In this paper, we propose a Density-Aware Feature Embedding Network (DA-Net) for the task of face clustering, which utilizes both local and non-local information, to learn a robust feature embedding.

664, TITLE: Self-Supervised Learning of Pretext-Invariant Representations

http://openaccess.thecvf.com/content_CVPR_2020/html/Misra_Self-Supervised_Learning_of_Pretext-Invariant_Representations_CVPR_2020_paper.html
AUTHORS: Ishan Misra, Laurens van der Maaten
HIGHLIGHT: Specifically, we develop Pretext-Invariant Representation Learning (PIRL, pronounced as `pearl') that learns invariant representations based on pretext tasks.

665, TITLE: ROAM: Recurrently Optimizing Tracking Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_ROAM_Recurrently_Optimizing_Tracking_Model_CVPR_2020_paper.html
AUTHORS: Tianyu Yang, Pengfei Xu, Runbo Hu, Hua Chai, Antoni B. Chan
HIGHLIGHT: In this paper, we design a tracking model consisting of response generation and bounding box regression, where the first component produces a heat map to indicate the presence of the object at different positions and the second part regresses the relative bounding box shifts to anchors mounted on sliding-window locations.

666, TITLE: Deformable Siamese Attention Networks for Visual Object Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Deformable_Siamese_Attention_Networks_for_Visual_Object_Tracking_CVPR_2020_paper.html
AUTHORS: Yuechen Yu, Yilei Xiong, Weilin Huang, Matthew R. Scott
HIGHLIGHT: In this paper, we propose Deformable Siamese Attention Networks, referred to as SiamAttn, by introducing a new Siamese attention mechanism that computes deformable self-attention and cross-attention.

667, TITLE: 15 Keypoints Is All You Need
http://openaccess.thecvf.com/content_CVPR_2020/html/Snower_15_Keypoints_Is_All_You_Need_CVPR_2020_paper.html
AUTHORS: Michael Snower, Asim Kadav, Farley Lai, Hans Peter Graf
HIGHLIGHT: We present an efficient multi-person pose-tracking method, KeyTrack that only relies on keypoint information without using any RGB or optical flow to locate and track human keypoints in real-time.

668, TITLE: Optical Flow in the Dark
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Optical_Flow_in_the_Dark_CVPR_2020_paper.html
AUTHORS: Yinqiang Zheng, Mingfang Zhang, Feng Lu
HIGHLIGHT: We propose an end-to-end data-driven method that avoids error accumulation and learns optical flow directly from low-light noisy images.
We also collect a new optical flow dataset in raw format with a large range of exposure to be used as a benchmark.

669, TITLE: Sketch-BERT: Learning Sketch Bidirectional Encoder Representation From Transformers by Self-Supervised Learning of Sketch Gestalt
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Sketch-BERT_Learning_Sketch_Bidirectional_Encoder_Representation_From_Transformers_by_Self-Supervised_CVPR_2020_paper.html
AUTHORS: Hangyu Lin, Yanwei Fu, Xiangyang Xue, Yu-Gang Jiang
HIGHLIGHT: Particularly, towards the pre-training task, we present a novel Sketch Gestalt Model (SGM) to help train the Sketch-BERT.

670, TITLE: A Unified Object Motion and Affinity Model for Online Multi-Object Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Yin_A_Unified_Object_Motion_and_Affinity_Model_for_Online_Multi-Object_CVPR_2020_paper.html
AUTHORS: Junbo Yin, Wenguan Wang, Qinghao Meng, Ruigang Yang, Jianbing Shen
HIGHLIGHT: In this paper, we propose a novel MOT framework that unifies object motion and affinity model into a single network, named UMA, in order to learn a compact feature that is discriminative for both object motion and affinity measure.

671, TITLE: Sub-Frame Appearance and 6D Pose Estimation of Fast Moving Objects
http://openaccess.thecvf.com/content_CVPR_2020/html/Rozumnyi_Sub-Frame_Appearance_and_6D_Pose_Estimation_of_Fast_Moving_Objects_CVPR_2020_paper.html
AUTHORS: Denys Rozumnyi, Jan Kotera, Filip Sroubek, Jiri Matas
HIGHLIGHT: We propose a novel method that tracks fast moving objects, mainly non-uniform spherical, in full 6 degrees of freedom, estimating simultaneously their 3D motion trajectory, 3D pose and object appearance changes with a time step that is a fraction of the video frame exposure time.

672, TITLE: How to Train Your Deep Multi-Object Tracker
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_How_to_Train_Your_Deep_Multi-Object_Tracker_CVPR_2020_paper.html
AUTHORS: Yihong Xu, Aljosa Osep, Yutong Ban, Radu Horaud, Laura Leal-Taixe, Xavier Alameda-Pineda

HIGHLIGHT: In this paper, we bridge this gap by proposing a differentiable proxy of MOTA and MOTP, which we combine in a loss function suitable for end-to-end training of deep multi-object trackers.

673, TITLE: TPNet: Trajectory Proposal Network for Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Fang_TPNet_Trajectory_Proposal_Network_for_Motion_Prediction_CVPR_2020_paper.html
AUTHORS: Liangji Fang, Qinhong Jiang, Jianping Shi, Bolei Zhou
HIGHLIGHT: In this work we propose a novel two-stage motion prediction framework, Trajectory Proposal Network (TPNet).

674, TITLE: Large Scale Video Representation Learning via Relational Graph Clustering
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Large_Scale_Video_Representation_Learning_via_Relational_Graph_Clustering_CVPR_2020_paper.html
AUTHORS: Hyodong Lee, Joonseok Lee, Joe Yue-Hei Ng, Paul Natsev
HIGHLIGHT: In this work, we explore two promising scalable representation learning approaches on video domain.

675, TITLE: Towards Universal Representation Learning for Deep Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Towards_Universal_Representation_Learning_for_Deep_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Yichun Shi, Xiang Yu, Kihyuk Sohn, Manmohan Chandraker, Anil K. Jain
HIGHLIGHT: Instead, we propose a universal representation learning framework that can deal with larger variation unseen in the given training data without leveraging target domain knowledge.

676, TITLE: Robust Partial Matching for Person Search in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhong_Robust_Partial_Matching_for_Person_Search_in_the_Wild_CVPR_2020_paper.html
AUTHORS: Yingji Zhong, Xiaoyu Wang, Shiliang Zhang
HIGHLIGHT: To alleviate this issue, this paper proposes an Align-to-Part Network (APNet) for person detection and re-Identification (reID).

677, TITLE: Correlation-Guided Attention for Corner Detection Based Visual Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Du_Correlation-Guided_Attention_for_Corner_Detection_Based_Visual_Tracking_CVPR_2020_paper.html
AUTHORS: Fei Du, Peng Liu, Wei Zhao, Xianglong Tang
HIGHLIGHT: We analyze the reasons for their failure and propose a state-of-the-art tracker that performs correlation-guided attentional corner detection in two stages.

678, TITLE: Learning Multi-Object Tracking and Segmentation From Automatic Annotations
http://openaccess.thecvf.com/content_CVPR_2020/html/Porzi_Learning_Multi-Object_Tracking_and_Segmentation_From_Automatic_Annotations_CVPR_2020_paper.html
AUTHORS: Lorenzo Porzi, Markus Hofinger, Idoia Ruiz, Joan Serrat, Samuel Rota Bulo, Peter Kontschieder
HIGHLIGHT: In this work we contribute a novel pipeline to automatically generate training data, and to improve over state-of-the-art multi-object tracking and segmentation (MOTS) methods.

679, TITLE: PandaNet: Anchor-Based Single-Shot Multi-Person 3D Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Benzine_PandaNet_Anchor-Based_Single-Shot_Multi-Person_3D_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Abdallah Benzine, Florian Chabot, Bertrand Luvison, Quoc Cuong Pham, Catherine Achard
HIGHLIGHT: In this work, we present PandaNet (Pose estimAtioN and Dectection Anchor-based Network), a new single-shot, anchor-based and multi-person 3D pose estimation approach.

680, TITLE: Rotation Consistent Margin Loss for Efficient Low-Bit Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Rotation_Consistent_Margin_Loss_for_Efficient_Low-Bit_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Yudong Wu, Yichao Wu, Ruihao Gong, Yuanhao Lv, Ken Chen, Ding Liang, Xiaolin Hu, Xianglong Liu, Junjie Yan
HIGHLIGHT: In this paper, we consider the low-bit quantization problem of face recognition (FR) under the open-set protocol.

681, TITLE: Joint Spatial-Temporal Optimization for Stereo 3D Object Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Joint_Spatial-Temporal_Optimization_for_Stereo_3D_Object_Tracking_CVPR_2020_paper.html

AUTHORS: Peiliang Li, Jieqi Shi, Shaojie Shen
HIGHLIGHT: To benefit from both the powerful object understanding skill from deep neural network meanwhile tackle precise geometry modeling for consistent trajectory estimation, we propose a joint spatial-temporal optimization-based stereo 3D object tracking method.

682, TITLE: Unity Style Transfer for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Unity_Style_Transfer_for_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Chong Liu, Xiaojun Chang, Yi-Dong Shen
HIGHLIGHT: To solve this problem, we propose a UnityStyle adaption method, which can smooth the style disparities within the same camera and across different cameras.

683, TITLE: Suppressing Uncertainties for Large-Scale Facial Expression Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Suppressing_Uncertainties_for_Large-Scale_Facial_Expression_Recognition_CVPR_2020_paper.html
AUTHORS: Kai Wang, Xiaojiang Peng, Jianfei Yang, Shijian Lu, Yu Qiao
HIGHLIGHT: To address this problelm, this paper proposes to suppress the uncertainties by a simple yet efficient Self-Cure Network (SCN).

684, TITLE: Multiview-Consistent Semi-Supervised Learning for 3D Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Mitra_Multiview-Consistent_Semi-Supervised_Learning_for_3D_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Rahul Mitra, Nitesh B. Gundavarapu, Abhishek Sharma, Arjun Jain
HIGHLIGHT: To reduce this annotation dependency, we propose Multiview-Consistent Semi Supervised Learning (MCSS) framework that utilizes similarity in pose information from unannotated, uncalibrated but synchronized multi-view videos of human motions as additional weak supervision signal to guide 3D human pose regression.

685, TITLE: Regularizing Neural Networks via Minimizing Hyperspherical Energy
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Regularizing_Neural_Networks_via_Minimizing_Hyperspherical_Energy_CVPR_2020_paper.html
AUTHORS: Rongmei Lin, Weiyang Liu, Zhen Liu, Chen Feng, Zhiding Yu, James M. Rehg, Li Xiong, Le Song
HIGHLIGHT: To address these problems, we propose the compressive minimum hyperspherical energy (CoMHE) as a more effective regularization for neural networks.

686, TITLE: Learning Representations by Predicting Bags of Visual Words
http://openaccess.thecvf.com/content_CVPR_2020/html/Gidaris_Learning_Representations_by_Predicting_Bags_of_Visual_Words_CVPR_2020_paper.html
AUTHORS: Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Perez, Matthieu Cord
HIGHLIGHT: Inspired by the success of NLP methods in this area, in this work we propose a self-supervised approach based on spatially dense image descriptions that encode discrete visual concepts, here called visual words.

687, TITLE: AnimalWeb: A Large-Scale Hierarchical Dataset of Annotated Animal Faces
http://openaccess.thecvf.com/content_CVPR_2020/html/Khan_AnimalWeb_A_Large-Scale_Hierarchical_Dataset_of_Annotated_Animal_Faces_CVPR_2020_paper.html
AUTHORS: Muhammad Haris Khan, John McDonagh, Salman Khan, Muhammad Shahabuddin, Aditya Arora, Fahad Shahbaz Khan, Ling Shao, Georgios Tzimiropoulos
HIGHLIGHT: To this end, we introduce a large-scale, hierarchical annotated dataset of animal faces, featuring 22.4K faces from 350 diverse species and 21 animal orders across biological taxonomy.

688, TITLE: A Transductive Approach for Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_A_Transductive_Approach_for_Video_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Yizhuo Zhang, Zhirong Wu, Houwen Peng, Stephen Lin
HIGHLIGHT: To address this issue, we propose a simple yet strong transductive method, in which additional modules, datasets, and dedicated architectural designs are not needed.

689, TITLE: Dynamic Face Video Segmentation via Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Dynamic_Face_Video_Segmentation_via_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS: Yujiang Wang, Mingzhi Dong, Jie Shen, Yang Wu, Shiyang Cheng, Maja Pantic

HIGHLIGHT: To overcome this limitation, we model the online key decision process in dynamic video segmentation as a deep reinforcement learning problem and learn an efficient and effective scheduling policy from expert information about decision history and from the process of maximising global return.

690, TITLE: Implicit Functions in Feature Space for 3D Shape Reconstruction and Completion
http://openaccess.thecvf.com/content_CVPR_2020/html/Chibane_Implicit_Functions_in_Feature_Space_for_3D_Shape_Reconstruction_and_CVPR_2020_paper.html
AUTHORS: Julian Chibane, Thiemo Alldieck, Gerard Pons-Moll
HIGHLIGHT: To solve this, we propose Implicit Feature Networks (IF-Nets), which deliver continuous outputs, can handle multiple topologies, and complete shapes for missing or sparse input data retaining the nice properties of recent learned implicit functions, but critically they can also retain detail when it is present in the input data, and can reconstruct articulated humans.

691, TITLE: Semantic Drift Compensation for Class-Incremental Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Semantic_Drift_Compensation_for_Class-Incremental_Learning_CVPR_2020_paper.html
AUTHORS: Lu Yu, Bartlomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, Joost van de Weijer
HIGHLIGHT: Therefore, we study incremental learning for embedding networks. In addition, we propose a new method to estimate the drift, called semantic drift, of features and compensate for it without the need of any exemplars.

692, TITLE: Context-Aware Human Motion Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Corona_Context-Aware_Human_Motion_Prediction_CVPR_2020_paper.html
AUTHORS: Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer
HIGHLIGHT: In this paper, we explore this scenario using a novel context-aware motion prediction architecture.

693, TITLE: DeepDeform: Learning Non-Rigid RGB-D Reconstruction With Semi-Supervised Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Bozic_DeepDeform_Learning_Non-Rigid_RGB-D_Reconstruction_With_Semi-Supervised_Data_CVPR_2020_paper.html
AUTHORS: Aljaz Bozic, Michael Zollhofer, Christian Theobalt, Matthias Niessner
HIGHLIGHT: Based on this corpus, we introduce a data-driven non-rigid feature matching approach, which we integrate into an optimization-based reconstruction pipeline.

694, TITLE: Optical Non-Line-of-Sight Physics-Based 3D Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Isogawa_Optical_Non-Line-of-Sight_Physics-Based_3D_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Mariko Isogawa, Ye Yuan, Matthew O'Toole, Kris M. Kitani
HIGHLIGHT: We describe a method for 3D human pose estimation from transient images (i.e., a 3D spatio-temporal histogram of photons) acquired by an optical non-line-of-sight (NLOS) imaging system.

695, TITLE: Learning to Transfer Texture From Clothing Images to 3D Humans
http://openaccess.thecvf.com/content_CVPR_2020/html/Mir_Learning_to_Transfer_Texture_From_Clothing_Images_to_3D_Humans_CVPR_2020_paper.html
AUTHORS: Aymen Mir, Thiemo Alldieck, Gerard Pons-Moll
HIGHLIGHT: In this paper, we present a simple yet effective method to automatically transfer textures of clothing images (front and back) to 3D garments worn on top SMPL, in real time.

696, TITLE: UniPose: Unified Human Pose Estimation in Single Images and Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Artacho_UniPose_Unified_Human_Pose_Estimation_in_Single_Images_and_Videos_CVPR_2020_paper.html
AUTHORS: Bruno Artacho, Andreas Savakis
HIGHLIGHT: We propose UniPose, a unified framework for human pose estimation, based on our "Waterfall" Atrous Spatial Pooling architecture, that achieves state-of-art-results on several pose estimation metrics.

697, TITLE: Minimal Solutions to Relative Pose Estimation From Two Views Sharing a Common Direction With Unknown Focal Length
http://openaccess.thecvf.com/content_CVPR_2020/html/Ding_Minimal_Solutions_to_Relative_Pose_Estimation_From_Two_Views_Sharing_CVPR_2020_paper.html
AUTHORS: Yaqing Ding, Jian Yang, Jean Ponce, Hui Kong
HIGHLIGHT: We propose minimal solutions to relative pose estimation problem from two views sharing a common direction with unknown focal length.

698, TITLE: 3D Human Mesh Regression With Dense Correspondence

http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_3D_Human_Mesh_Regression_With_Dense_Correspondence_CVPR_2020_paper.html

AUTHORS: Wang Zeng, Wanli Ouyang, Ping Luo, Wentao Liu, Xiaogang Wang

HIGHLIGHT: This paper proposes a model-free 3D human mesh estimation framework, named DecoMR, which explicitly establishes the dense correspondence between the mesh and the local image features in the UV space (i.e. a 2D space used for texture mapping of 3D mesh).

699, TITLE: Cross-Modal Pattern-Propagation for RGB-T Tracking

http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Cross-Modal_Pattern-Propagation_for_RGB-T_Tracking_CVPR_2020_paper.html

AUTHORS: Chaoqun Wang, Chunyan Xu, Zhen Cui, Ling Zhou, Tong Zhang, Xiaoya Zhang, Jian Yang

HIGHLIGHT: Motivated by our observations on RGB-T data that pattern correlations are high-frequently recurred across modalities also along sequence frames, in this paper, we propose a cross-modal pattern-propagation (CMPP) tracking framework to diffuse instance patterns across RGB-T data on spatial domain as well as temporal domain.

700, TITLE: Distilling Knowledge From Graph Convolutional Networks

http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Distilling_Knowledge_From_Graph_Convolutional_Networks_CVPR_2020_paper.html

AUTHORS: Yiding Yang, Jiayan Qiu, Mingli Song, Dacheng Tao, Xinchao Wang

HIGHLIGHT: In this paper, we propose to our best knowledge the first dedicated approach to distilling knowledge from a pre-trained GCN model.

701, TITLE: Learning Identity-Invariant Motion Representations for Cross-ID Face Reenactment

http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Learning_Identity-Invariant_Motion_Representations_for_Cross-ID_Face_Reenactment_CVPR_2020_paper.html

AUTHORS: Po-Hsiang Huang, Fu-En Yang, Yu-Chiang Frank Wang

HIGHLIGHT: In this paper, we propose a unique network of CrossID-GAN to perform multi-ID face reenactment.

702, TITLE: Distribution-Aware Coordinate Representation for Human Pose Estimation

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Distribution-Aware_Coordinate_Representation_for_Human_Pose_Estimation_CVPR_2020_paper.html

AUTHORS: Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, Ce Zhu

HIGHLIGHT: For the first time, we find that the process of decoding the predicted heatmaps into the final joint coordinates in the original image space is surprisingly significant for the performance. We further probe the design limitations of the standard coordinate decoding method, and propose a more principled distributionaware decoding method.

703, TITLE: Parsing-Based View-Aware Embedding Network for Vehicle Re-Identification

http://openaccess.thecvf.com/content_CVPR_2020/html/Meng_Parsing-Based_View-Aware_Embedding_Network_for_Vehicle_Re-Identification_CVPR_2020_paper.html

AUTHORS: Dechao Meng, Liang Li, Xuejing Liu, Yadong Li, Shijie Yang, Zheng-Jun Zha, Xingyu Gao, Shuhui Wang, Qingming Huang

HIGHLIGHT: In this paper, we propose a parsing-based view-aware embedding network (PVEN) to achieve the view-aware feature alignment and enhancement for vehicle ReID.

704, TITLE: HandVoxNet: Deep Voxel-Based Network for 3D Hand Shape and Pose Estimation From a Single Depth Map

http://openaccess.thecvf.com/content_CVPR_2020/html/Malik_HandVoxNet_Deep_Voxel-Based_Network_for_3D_Hand_Shape_and_Pose_CVPR_2020_paper.html

AUTHORS: Jameel Malik, Ibrahim Abdelaziz, Ahmed Elhayek, Soshi Shimada, Sk Aziz Ali, Vladislav Golyanik, Christian Theobalt, Didier Stricker

HIGHLIGHT: In contrast, we propose a novel architecture with 3D convolutions trained in a weakly-supervised manner.

705, TITLE: Determinant Regularization for Gradient-Efficient Graph Matching

http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Determinant_Regularization_for_Gradient-Efficient_Graph_Matching_CVPR_2020_paper.html

AUTHORS: Tianshu Yu, Junchi Yan, Baoxin Li

HIGHLIGHT: In this paper, we show a novel regularization technique with the tool of determinant analysis on the matching matrix which is relaxed into continuous domain with gradient based optimization.

706, TITLE: D3S - A Discriminative Single Shot Segmentation Tracker

http://openaccess.thecvf.com/content_CVPR_2020/html/Lukezic_D3S_-
_A_Discriminative_Single_Shot_Segmentation_Tracker_CVPR_2020_paper.html
AUTHORS: Alan Lukezic, Jiri Matas, Matej Kristan
HIGHLIGHT: We propose a discriminative single-shot segmentation tracker - D3S, which narrows the gap between visual object tracking and video object segmentation.

707, TITLE: MANTRA: Memory Augmented Networks for Multiple Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Marchetti_MANTRA_Memory_Augmented_Networks_for_Multiple_Trajectory_Prediction_CVPR_2020_paper.html
AUTHORS: Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, Alberto Del Bimbo
HIGHLIGHT: In this paper we address the problem of multimodal trajectory prediction exploiting a Memory Augmented Neural Network.

708, TITLE: End-to-End Model-Free Reinforcement Learning for Urban Driving Using Implicit Affordances
http://openaccess.thecvf.com/content_CVPR_2020/html/Toromanoff_End-to-End_Model-
Free_Reinforcement_Learning_for_Urban_Driving_Using_Implicit_Affordances_CVPR_2020_paper.html
AUTHORS: Marin Toromanoff, Emilie Wirbel, Fabien Moutarde
HIGHLIGHT: We present a novel technique, coined implicit affordances, to effectively leverage RL for urban driving thus including lane keeping, pedestrians and vehicles avoidance, and traffic light detection.

709, TITLE: GraphTER: Unsupervised Learning of Graph Transformation Equivariant Representations via Auto-Encoding Node-Wise Transformations
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_GraphTER_Unsupervised_Learning_of_Graph_Transformation_Equivariant_Representations_via_Auto-Encoding_CVPR_2020_paper.html
AUTHORS: Xiang Gao, Wei Hu, Guo-Jun Qi
HIGHLIGHT: To this end, we propose a novel unsupervised learning of Graph Transformation Equivariant Representations (GraphTER), aiming to capture intrinsic patterns of graph structure under both global and local transformations.

710, TITLE: Can Facial Pose and Expression Be Separated With Weak Perspective Camera?
http://openaccess.thecvf.com/content_CVPR_2020/html/Sariyanidi_Can_Facial_Pose_and_Expression_Be_Separated_With_Weak_Perspective_CVPR_2020_paper.html
AUTHORS: Evangelos Sariyanidi, Casey J. Zampella, Robert T. Schultz, Birkan Tunc
HIGHLIGHT: This paper critically examines the suitability of WP camera for separating facial pose and expression.

711, TITLE: Probabilistic Regression for Visual Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Danelljan_Probabilistic_Regression_for_Visual_Tracking_CVPR_2020_paper.html
AUTHORS: Martin Danelljan, Luc Van Gool, Radu Timofte
HIGHLIGHT: In this work, we therefore propose a probabilistic regression formulation and apply it to tracking.

712, TITLE: 3DRegNet: A Deep Neural Network for 3D Point Registration
http://openaccess.thecvf.com/content_CVPR_2020/html/Pais_3DRegNet_A_Deep_Neural_Network_for_3D_Point_Registration_CVPR_2020_paper.html
AUTHORS: G. Dias Pais, Srikumar Ramalingam, Venu Madhav Govindu, Jacinto C. Nascimento, Rama Chellappa, Pedro Miraldo
HIGHLIGHT: We present 3DRegNet, a novel deep learning architecture for the registration of 3D scans.

713, TITLE: Compressed Volumetric Heatmaps for Multi-Person 3D Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Fabbri_Compressed_Volumetric_Heatmaps_for_Multi-
Person_3D_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Matteo Fabbri, Fabio Lanzi, Simone Calderara, Stefano Alletto, Rita Cucchiara
HIGHLIGHT: In this paper we present a novel approach for bottom-up multi-person 3D human pose estimation from monocular RGB images.

714, TITLE: Three-Dimensional Reconstruction of Human Interactions
http://openaccess.thecvf.com/content_CVPR_2020/html/Fieraru_Three-
Dimensional_Reconstruction_of_Human_Interactions_CVPR_2020_paper.html
AUTHORS: Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, Cristian Sminchisescu
HIGHLIGHT: This paper addresses such issues and makes several contributions: (1) we introduce models for interaction signature estimation (ISP) encompassing contact detection, segmentation, and 3d contact signature prediction; (2) we show how such components can be leveraged in order to produce augmented losses that ensure contact consistency during 3d reconstruction; (3) we construct several large datasets for learning and evaluating 3d contact prediction and reconstruction methods; specifically, we

introduce CHI3D, a lab-based accurate 3d motion capture dataset with 631 sequences containing 2,525 contact events, 728,664 ground truth 3d poses, as well as FlickrCI3D, a dataset of 11,216 images, with 14,081 processed pairs of people, and 81,233 facet-level surface correspondences within 138,213 selected contact regions.

715, TITLE:	Distribution-Induced Bidirectional Generative Adversarial Network for Graph Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Distribution-Induced_Bidirectional_Generative_Adversarial_Network_for_Graph_Representation_Learning_CVPR_2020_paper.html
AUTHORS:	Shuai Zheng, Zhenfeng Zhu, Xingxing Zhang, Zhizhe Liu, Jian Cheng, Yao Zhao
HIGHLIGHT:	In this paper, we propose a Distribution-induced Bidirectional Generative Adversarial Network (named DBGAN) for graph representation learning.

716, TITLE:	Minimal Solvers for 3D Scan Alignment With Pairs of Intersecting Lines
http://openaccess.thecvf.com/content_CVPR_2020/html/Mateus_Minimal_Solvers_for_3D_Scan_Alignment_With_Pairs_of_Intersecting_CVPR_2020_paper.html
AUTHORS:	Andre Mateus, Srikumar Ramalingam, Pedro Miraldo
HIGHLIGHT:	In this paper, we present minimal solvers that combine these different type of constraints: 1) three line intersections and one point match; 2) one line intersection and two point matches; 3) three line intersections and one plane match; 4) one line intersection and two plane matches; and 5) one line intersection, one point match, and one plane match.

717, TITLE:	Wavelet Integrated CNNs for Noise-Robust Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Wavelet_Integrated_CNNs_for_Noise-Robust_Image_Classification_CVPR_2020_paper.html
AUTHORS:	Qiufu Li, Linlin Shen, Sheng Guo, Zhihui Lai
HIGHLIGHT:	We present general DWT and Inverse DWT (IDWT) layers applicable to various wavelets like Haar, Daubechies, and Cohen, etc., and design wavelet integrated CNNs (WaveCNets) using these layers for image classification.

718, TITLE:	Embedding Expansion: Augmentation in Embedding Space for Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Ko_Embedding_Expansion_Augmentation_in_Embedding_Space_for_Deep_Metric_Learning_CVPR_2020_paper.html
AUTHORS:	Byungsoo Ko, Geonmo Gu
HIGHLIGHT:	In this paper, inspired by query expansion and database augmentation, we propose an augmentation method in an embedding space for pair-based metric learning losses, called embedding expansion.

719, TITLE:	PropagationNet: Propagate Points to Curve to Learn Structure Information
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_PropagationNet_Propagate_Points_to_Curve_to_Learn_Structure_Information_CVPR_2020_paper.html
AUTHORS:	Xiehe Huang, Weihong Deng, Haifeng Shen, Xiubao Zhang, Jieping Ye
HIGHLIGHT:	In this paper, we explore the instincts and reasons behind our two proposals, i.e. Propagation Module and Focal Wing Loss, to tackle the problem.

720, TITLE:	Sequential 3D Human Pose and Shape Estimation From Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Sequential_3D_Human_Pose_and_Shape_Estimation_From_Point_Clouds_CVPR_2020_paper.html
AUTHORS:	Kangkan Wang, Jin Xie, Guofeng Zhang, Lei Liu, Jian Yang
HIGHLIGHT:	In this paper, we propose a novel sequential 3D human pose and shape estimation framework from a sequence of point clouds.

721, TITLE:	Improving the Robustness of Capsule Networks to Image Affine Transformations
http://openaccess.thecvf.com/content_CVPR_2020/html/Gu_Improving_the_Robustness_of_Capsule_Networks_to_Image_Affine_Transformations_CVPR_2020_paper.html
AUTHORS:	Jindong Gu, Volker Tresp
HIGHLIGHT:	Furthermore, we explore the limitations of capsule transformations and propose affine CapsNets (Aff-CapsNets), which are more robust to affine transformations.

722, TITLE:	Noise Modeling, Synthesis and Classification for Generic Object Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2020/html/Stehouwer_Noise_Modeling_Synthesis_and_Classification_for_Generic_Object_Anti-Spoofing_CVPR_2020_paper.html
AUTHORS:	Joel Stehouwer, Amin Jourabloo, Yaojie Liu, Xiaoming Liu
HIGHLIGHT:	In this work, we define and tackle the problem of Generic Object Anti-Spoofing (GOAS) for the first time.

723, TITLE:	Quaternion Product Units for Deep Learning on 3D Rotation Groups

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Quaternion_Product_Units_for_Deep_Learning_on_3D_Rotation_Groups_CVPR_2020_paper.html
AUTHORS: Xuan Zhang, Shaofei Qin, Yi Xu, Hongteng Xu
HIGHLIGHT: We propose a novel quaternion product unit (QPU) to represent data on 3D rotation groups.

724, TITLE: Unsupervised Representation Learning for Gaze Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Unsupervised_Representation_Learning_for_Gaze_Estimation_CVPR_2020_paper.html
AUTHORS: Yu Yu, Jean-Marc Odobez
HIGHLIGHT: To address this issue, our main contribution in this paper is to propose an effective approach to learn a low dimensional gaze representation without gaze annotations, which to the best of our best knowledge, is the first work to do so.

725, TITLE: P-nets: Deep Polynomial Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Chrysos_P-nets_Deep_Polynomial_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Grigorios G. Chrysos, Stylianos Moschoglou, Giorgos Bouritsas, Yannis Panagakis, Jiankang Deng, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we propose \Pi-Nets, a new class of DCNNs.

726, TITLE: Hierarchically Robust Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Qian_Hierarchically_Robust_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Qi Qian, Juhua Hu, Hao Li
HIGHLIGHT: In this work, we investigate this phenomenon and demonstrate that deep features can be suboptimal due to the fact that they are learned by minimizing the empirical risk.

727, TITLE: How Useful Is Self-Supervised Pretraining for Visual Tasks?
http://openaccess.thecvf.com/content_CVPR_2020/html/Newell_How_Useful_Is_Self-Supervised_Pretraining_for_Visual_Tasks_CVPR_2020_paper.html
AUTHORS: Alejandro Newell, Jia Deng
HIGHLIGHT: We investigate what factors may play a role in the utility of these pretraining methods for practitioners.

728, TITLE: Copy and Paste GAN: Face Hallucination From Shaded Thumbnails
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Copy_and_Paste_GAN_Face_Hallucination_From_Shaded_Thumbnails_CVPR_2020_paper.html
AUTHORS: Yang Zhang, Ivor W. Tsang, Yawei Luo, Chang-Hui Hu, Xiaobo Lu, Xin Yu
HIGHLIGHT: This paper proposes a Copy and Paste Generative Adversarial Network (CPGAN) to recover authentic high-resolution (HR) face images while compensating for low and non-uniform illumination.

729, TITLE: TailorNet: Predicting Clothing in 3D as a Function of Human Pose, Shape and Garment Style
http://openaccess.thecvf.com/content_CVPR_2020/html/Patel_TailorNet_Predicting_Clothing_in_3D_as_a_Function_of_Human_CVPR_2020_paper.html
AUTHORS: Chaitanya Patel, Zhouyingcheng Liao, Gerard Pons-Moll
HIGHLIGHT: In this paper, we present TailorNet, a neural model which predicts clothing deformation in 3D as a function of three factors: pose, shape and style (garment geometry), while retaining wrinkle detail.

730, TITLE: Object-Occluded Human Shape and Pose Estimation From a Single Color Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Object-Occluded_Human_Shape_and_Pose_Estimation_From_a_Single_Color_CVPR_2020_paper.html
AUTHORS: Tianshu Zhang, Buzhen Huang, Yangang Wang
HIGHLIGHT: In this paper, we focus on the problem of directly estimating the object-occluded human shape and pose from single color images.
To supervise the network training, we further build a novel dataset named as 3DOH50K.

731, TITLE: Recursive Least-Squares Estimator-Aided Online Learning for Visual Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Recursive_Least-Squares_Estimator-Aided_Online_Learning_for_Visual_Tracking_CVPR_2020_paper.html
AUTHORS: Jin Gao, Weiming Hu, Yan Lu
HIGHLIGHT: Despite many dedicated techniques proposed to somehow treat those issues, in this paper we take a new way to strike a compromise between them based on the recursive least-squares estimation (LSE) algorithm.

732, TITLE: Self-Supervised Monocular Scene Flow Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hur_Self-
Supervised_Monocular_Scene_Flow_Estimation_CVPR_2020_paper.html
AUTHORS: Junhwa Hur, Stefan Roth
HIGHLIGHT: We propose a novel monocular scene flow method that yields competitive accuracy and real-time performance.

733, TITLE: Learning Fast and Robust Target Models for Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Robinson_Learning_Fast_and_Robust_Target_Models_for_Video_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Andreas Robinson, Felix Jaremo Lawin, Martin Danelljan, Fahad Shahbaz Khan, Michael Felsberg
HIGHLIGHT: We propose a novel VOS architecture consisting of two network components.

734, TITLE: Reciprocal Learning Networks for Human Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Reciprocal_Learning_Networks_for_Human_Trajectory_Prediction_CVPR_2020_paper.html
AUTHORS: Hao Sun, Zhiqun Zhao, Zhihai He
HIGHLIGHT: Based on this unique property, we develop a new approach, called reciprocal learning, for human trajectory prediction.

735, TITLE: Nonparametric Object and Parts Modeling With Lie Group Dynamics
http://openaccess.thecvf.com/content_CVPR_2020/html/Hayden_Nonparametric_Object_and_Parts_Modeling_With_Lie_Group_Dynamics_CVPR_2020_paper.html
AUTHORS: David S. Hayden, Jason Pacheco, John W. Fisher III
HIGHLIGHT: Here, we relax such strong assumptions via an unsupervised, Bayesian nonparametric parts model that infers an unknown number of parts with motions coupled by a body dynamic and parameterized by SE(D), the Lie group of rigid transformations.

736, TITLE: Learning to Shadow Hand-Drawn Sketches
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Learning_to_Shadow_Hand-
Drawn_Sketches_CVPR_2020_paper.html
AUTHORS: Qingyuan Zheng, Zhuoru Li, Adam Bargteil
HIGHLIGHT: We present a fully automatic method to generate detailed and accurate artistic shadows from pairs of line drawing sketches and lighting directions.

737, TITLE: Intuitive, Interactive Beard and Hair Synthesis With Generative Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Olszewski_Intuitive_Interactive_Beard_and_Hair_Synthesis_With_Generative_Models_CVPR_2020_paper.html
AUTHORS: Kyle Olszewski, Duygu Ceylan, Jun Xing, Jose Echevarria, Zhili Chen, Weikai Chen, Hao Li
HIGHLIGHT: We present an interactive approach to synthesizing realistic variations in facial hair in images, ranging from subtle edits to existing hair to the addition of complex and challenging hair in images of clean-shaven subjects.

738, TITLE: Semantic Pyramid for Image Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Shocher_Semantic_Pyramid_for_Image_Generation_CVPR_2020_paper.html
AUTHORS: Assaf Shocher, Yossi Gandelsman, Inbar Mosseri, Michal Yarom, Michal Irani, William T. Freeman, Tali Dekel
HIGHLIGHT: We present a novel GAN-based model that utilizes the space of deep features learned by a pre-trained classification model.

739, TITLE: SynSin: End-to-End View Synthesis From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Wiles_SynSin_End-to-
End_View_Synthesis_From_a_Single_Image_CVPR_2020_paper.html
AUTHORS: Olivia Wiles, Georgia Gkioxari, Richard Szeliski, Justin Johnson
HIGHLIGHT: We propose a novel end-to-end model for this task using a single image at test time; it is trained on real images without any ground-truth 3D information.

740, TITLE: A Characteristic Function Approach to Deep Implicit Generative Modeling
http://openaccess.thecvf.com/content_CVPR_2020/html/Ansari_A_Characteristic_Function_Approach_to_Deep_Implicit_Generative_Modeling_CVPR_2020_paper.html
AUTHORS: Abdul Fatir Ansari, Jonathan Scarlett, Harold Soh
HIGHLIGHT: In this paper, we formulate the problem of learning an IGM as minimizing the expected distance between characteristic functions.

741, TITLE:	High-Resolution Daytime Translation Without Domain Labels
http://openaccess.thecvf.com/content_CVPR_2020/html/Anokhin_High-Resolution_Daytime_Translation_Without_Domain_Labels_CVPR_2020_paper.html
AUTHORS:	Ivan Anokhin, Pavel Solovev, Denis Korzhenkov, Alexey Kharlamov, Taras Khakhulin, Aleksei Silvestrov, Sergey Nikolenko, Victor Lempitsky, Gleb Sterkin
HIGHLIGHT:	We present the high-resolution daytime translation (HiDT) model for this task.

742, TITLE:	Leveraging 2D Data to Learn Textured 3D Mesh Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Henderson_Leveraging_2D_Data_to_Learn_Textured_3D_Mesh_Generation_CVPR_2020_paper.html
AUTHORS:	Paul Henderson, Vagia Tsiminaki, Christoph H. Lampert
HIGHLIGHT:	In this work, we present the first generative model of textured 3D meshes.

743, TITLE:	Contextual Residual Aggregation for Ultra High-Resolution Image Inpainting
http://openaccess.thecvf.com/content_CVPR_2020/html/Yi_Contextual_Residual_Aggregation_for_Ultra_High-Resolution_Image_Inpainting_CVPR_2020_paper.html
AUTHORS:	Zili Yi, Qiang Tang, Shekoofeh Azizi, Daesik Jang, Zhan Xu
HIGHLIGHT:	Motivated by this, we propose a Contextual Residual Aggregation (CRA) mechanism that can produce high-frequency residuals for missing contents by weighted aggregating residuals from contextual patches, thus only requiring a low-resolution prediction from the network.

744, TITLE:	Flow Contrastive Estimation of Energy-Based Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Flow_Contrastive_Estimation_of_Energy-Based_Models_CVPR_2020_paper.html
AUTHORS:	Ruiqi Gao, Erik Nijkamp, Diederik P. Kingma, Zhen Xu, Andrew M. Dai, Ying Nian Wu
HIGHLIGHT:	This paper studies a training method to jointly estimate an energy-based model and a flow-based model, in which the two models are iteratively updated based on a shared adversarial value function.

745, TITLE:	Hardware-in-the-Loop End-to-End Optimization of Camera Image Processing Pipelines
http://openaccess.thecvf.com/content_CVPR_2020/html/Mosleh_Hardware-in-the-Loop_End-to-End_Optimization_of_Camera_Image_Processing_Pipelines_CVPR_2020_paper.html
AUTHORS:	Ali Mosleh, Avinash Sharma, Emmanuel Onzon, Fahim Mannan, Nicolas Robidoux, Felix Heide
HIGHLIGHT:	Departing from such approximations, we present a hardware-in-the-loop method that directly optimizes hardware image processing pipelines for end-to-end domain-specific losses by solving a nonlinear multi-objective optimization problem with a novel 0th-order stochastic solver directly interfaced with the hardware ISP.

746, TITLE:	Search to Distill: Pearls Are Everywhere but Not the Eyes
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Search_to_Distill_Pearls_Are_Everywhere_but_Not_the_Eyes_CVPR_2020_paper.html
AUTHORS:	Yu Liu, Xuhui Jia, Mingxing Tan, Raviteja Vemulapalli, Yukun Zhu, Bradley Green, Xiaogang Wang
HIGHLIGHT:	To achieve this, we present a new Architecture-aware Knowledge Distillation (AKD) approach that finds student models (pearls for the teacher) that are best for distilling the given teacher model.

747, TITLE:	Total Deep Variation for Linear Inverse Problems
http://openaccess.thecvf.com/content_CVPR_2020/html/Kobler_Total_Deep_Variation_for_Linear_Inverse_Problems_CVPR_2020_paper.html
AUTHORS:	Erich Kobler, Alexander Effland, Karl Kunisch, Thomas Pock
HIGHLIGHT:	In this paper, we propose a novel learnable general-purpose regularizer exploiting recent architectural design patterns from deep learning.

748, TITLE:	Relative Interior Rule in Block-Coordinate Descent
http://openaccess.thecvf.com/content_CVPR_2020/html/Werner_Relative_Interior_Rule_in_Block-Coordinate_Descent_CVPR_2020_paper.html
AUTHORS:	Tomas Werner, Daniel Prusa, Tomas Dlask
HIGHLIGHT:	Based on this observation, we develop a theoretical framework for block-coordinate descent applied to general convex problems.

749, TITLE:	Learning Combinatorial Solver for Graph Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Learning_Combinatorial_Solver_for_Graph_Matching_CVPR_2020_paper.html

AUTHORS: Tao Wang, He Liu, Yidong Li, Yi Jin, Xiaohui Hou, Haibin Ling
HIGHLIGHT: In this paper we propose a fully trainable framework for graph matching, in which learning of affinities and solving for combinatorial optimization are not explicitly separated as in many previous arts.

750, TITLE: SampleNet: Differentiable Point Cloud Sampling
http://openaccess.thecvf.com/content_CVPR_2020/html/Lang_SampleNet_Differentiable_Point_Cloud_Sampling_CVPR_2020_paper.html
AUTHORS: Itai Lang, Asaf Manor, Shai Avidan
HIGHLIGHT: We introduce a novel differentiable relaxation for point cloud sampling that approximates sampled points as a mixture of points in the primary input cloud.

751, TITLE: Can We Learn Heuristics for Graphical Model Inference Using Reinforcement Learning?
http://openaccess.thecvf.com/content_CVPR_2020/html/Messaoud_Can_We_Learn_Heuristics_for_Graphical_Model_Inference_Using_Reinforcement_CVPR_2020_paper.html
AUTHORS: Safa Messaoud, Maghav Kumar, Alexander G. Schwing
HIGHLIGHT: In this paper, we show that we can learn program heuristics, i.e., policies, for solving inference in higher order CRFs for the task of semantic segmentation, using reinforcement learning.

752, TITLE: Quasi-Newton Solver for Robust Non-Rigid Registration
http://openaccess.thecvf.com/content_CVPR_2020/html/Yao_Quasi-Newton_Solver_for_Robust_Non-Rigid_Registration_CVPR_2020_paper.html
AUTHORS: Yuxin Yao, Bailin Deng, Weiwei Xu, Juyong Zhang
HIGHLIGHT: In this paper, we propose a formulation for robust non-rigid registration based on a globally smooth robust estimator for data fitting and regularization, which can handle outliers and partial overlaps.

753, TITLE: Rethinking Class-Balanced Methods for Long-Tailed Visual Recognition From a Domain Adaptation Perspective
http://openaccess.thecvf.com/content_CVPR_2020/html/Jamal_Rethinking_Class-Balanced_Methods_for_Long-Tailed_Visual_Recognition_From_a_Domain_CVPR_2020_paper.html
AUTHORS: Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, Boqing Gong
HIGHLIGHT: To this end, we propose to augment the classic class-balanced learning by explicitly estimating the differences between the class-conditioned distributions with a meta-learning approach.

754, TITLE: Optimizing Rank-Based Metrics With Blackbox Differentiation
http://openaccess.thecvf.com/content_CVPR_2020/html/Rolinek_Optimizing_Rank-Based_Metrics_With_Blackbox_Differentiation_CVPR_2020_paper.html
AUTHORS: Michal Rolinek, Vit Musil, Anselm Paulus, Marin Vlastelica, Claudio Michaelis, Georg Martius
HIGHLIGHT: We present an efficient, theoretically sound, and general method for differentiating rank-based metrics with mini-batch gradient descent.

755, TITLE: DualSDF: Semantic Shape Manipulation Using a Two-Level Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hao_DualSDF_Semantic_Shape_Manipulation_Using_a_Two-Level_Representation_CVPR_2020_paper.html
AUTHORS: Zekun Hao, Hadar Averbuch-Elor, Noah Snavely, Serge Belongie
HIGHLIGHT: We propose DualSDF, a representation expressing shapes at two levels of granularity, one capturing fine details and the other representing an abstracted proxy shape using simple and semantically consistent shape primitives.

756, TITLE: Dynamic Hierarchical Mimicking Towards Consistent Optimization Objectives
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Dynamic_Hierarchical_Mimicking_Towards_Consistent_Optimization_Objectives_CVPR_2020_paper.html
AUTHORS: Duo Li, Qifeng Chen
HIGHLIGHT: Complementary to previous training strategies, we propose Dynamic Hierarchical Mimicking, a generic feature learning mechanism, to advance CNN training with enhanced generalization ability.

757, TITLE: Deep Homography Estimation for Dynamic Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Le_Deep_Homography_Estimation_for_Dynamic_Scenes_CVPR_2020_paper.html
AUTHORS: Hoang Le, Feng Liu, Shu Zhang, Aseem Agarwala
HIGHLIGHT: This paper investigates and discusses how to design and train a deep neural network that handles dynamic scenes.

758, TITLE: PF-Net: Point Fractal Network for 3D Point Cloud Completion
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_PF-Net_Point_Fractal_Network_for_3D_Point_Cloud_Completion_CVPR_2020_paper.html
AUTHORS: Zitian Huang, Yikuan Yu, Jiawen Xu, Feng Ni, Xinyi Le
HIGHLIGHT: In this paper, we propose a Point Fractal Network (PF-Net), a novel learning-based approach for precise and high-fidelity point cloud completion.

759, TITLE: On the Regularization Properties of Structured Dropout
http://openaccess.thecvf.com/content_CVPR_2020/html/Pal_On_the_Regularization_Properties_of_Structured_Dropout_CVPR_2020_paper.html
AUTHORS: Ambar Pal, Connor Lane, Rene Vidal, Benjamin D. Haeffele
HIGHLIGHT: In this work we show that for single hidden-layer linear networks, DropBlock induces spectral k-support norm regularization, and promotes solutions that are low-rank and have factors with equal norm.

760, TITLE: Learning Oracle Attention for High-Fidelity Face Completion
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Learning_Oracle_Attention_for_High-Fidelity_Face_Completion_CVPR_2020_paper.html
AUTHORS: Tong Zhou, Changxing Ding, Shaowen Lin, Xinchao Wang, Dacheng Tao
HIGHLIGHT: Accordingly, in this paper, we design a comprehensive framework for face completion based on the U-Net structure.

761, TITLE: Deep Image Spatial Transformation for Person Image Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Ren_Deep_Image_Spatial_Transformation_for_Person_Image_Generation_CVPR_2020_paper.html
AUTHORS: Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H. Li, Ge Li
HIGHLIGHT: In this paper, we propose a differentiable global-flow local-attention framework to reassemble the inputs at the feature level.

762, TITLE: Learning to Optimize on SPD Manifolds
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Learning_to_Optimize_on_SPD_Manifolds_CVPR_2020_paper.html
AUTHORS: Zhi Gao, Yuwei Wu, Yunde Jia, Mehrtash Harandi
HIGHLIGHT: In this paper, we propose a meta-learning method to automatically learn an iterative optimizer on SPD manifolds.

763, TITLE: Deep 3D Portrait From a Single Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Deep_3D_Portrait_From_a_Single_Image_CVPR_2020_paper.html
AUTHORS: Sicheng Xu, Jiaolong Yang, Dong Chen, Fang Wen, Yu Deng, Yunde Jia, Xin Tong
HIGHLIGHT: In this paper, we present a learning-based approach for recovering the 3D geometry of human head from a single portrait image.

764, TITLE: RDCFace: Radial Distortion Correction for Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_RDCFace_Radial_Distortion_Correction_for_Face_Recognition_CVPR_2020_paper.html
AUTHORS: He Zhao, Xianghua Ying, Yongjie Shi, Xin Tong, Jingsi Wen, Hongbin Zha
HIGHLIGHT: In this paper, we propose a distortion-invariant face recognition system called RDCFace, which directly and only utilize the distorted images of faces, to alleviate the effects of radial lens distortion.

765, TITLE: Global-Local GCN: Large-Scale Label Noise Cleansing for Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Global-Local_GCN_Large-Scale_Label_Noise_Cleansing_for_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Yaobin Zhang, Weihong Deng, Mei Wang, Jiani Hu, Xian Li, Dongyue Zhao, Dongchao Wen
HIGHLIGHT: To solve this problem, we propose an effective automatic label noise cleansing framework for face recognition datasets, FaceGraph.

766, TITLE: MISC: Multi-Condition Injection and Spatially-Adaptive Compositing for Conditional Person Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Weng_MISC_Multi-Condition_Injection_and_Spatially-Adaptive_Compositing_for_Conditional_Person_Image_CVPR_2020_paper.html
AUTHORS: Shuchen Weng, Wenbo Li, Dawei Li, Hongxia Jin, Boxin Shi
HIGHLIGHT: In this paper, we explore synthesizing person images with multiple conditions for various backgrounds.

767, TITLE: SAINT: Spatially Aware Interpolation NeTwork for Medical Slice Synthesis

http://openaccess.thecvf.com/content_CVPR_2020/html/Peng_SAINT_Spatially_Aware_Interpolation_NeTwork_for_Medical_Slice_Synthesis_CVPR_2020_paper.html
AUTHORS: Cheng Peng, Wei-An Lin, Haofu Liao, Rama Chellappa, S. Kevin Zhou
HIGHLIGHT: In this paper, we introduce a Spatially Aware Interpolation NeTwork (SAINT) for medical slice synthesis to alleviate the memory constraint that volumetric data poses.

768, TITLE: Recurrent Feature Reasoning for Image Inpainting
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Recurrent_Feature_Reasoning_for_Image_Inpainting_CVPR_2020_paper.html
AUTHORS: Jingyuan Li, Ning Wang, Lefei Zhang, Bo Du, Dacheng Tao
HIGHLIGHT: In this paper, we devise a Recurrent Feature Reasoning (RFR) network which is mainly constructed by a plug-and-play Recurrent Feature Reasoning module and a Knowledge Consistent Attention (KCA) module.

769, TITLE: Structure-Preserving Super Resolution With Gradient Guidance
http://openaccess.thecvf.com/content_CVPR_2020/html/Ma_Structure-Preserving_Super_Resolution_With_Gradient_Guidance_CVPR_2020_paper.html
AUTHORS: Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, Jie Zhou
HIGHLIGHT: In this paper, we propose a structure-preserving super resolution method to alleviate the above issue while maintaining the merits of GAN-based methods to generate perceptual-pleasant details.

770, TITLE: Epipolar Transformers
http://openaccess.thecvf.com/content_CVPR_2020/html/He_Epipolar_Transformers_CVPR_2020_paper.html
AUTHORS: Yihui He, Rui Yan, Katerina Fragkiadaki, Shoou-I Yu
HIGHLIGHT: Therefore, we propose the differentiable "epipolar transformer", which enables the 2D detector to leverage 3D-aware features to improve 2D pose estimation.

771, TITLE: Diversified Arbitrary Style Transfer via Deep Feature Perturbation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Diversified_Arbitrary_Style_Transfer_via_Deep_Feature_Perturbation_CVPR_2020_paper.html
AUTHORS: Zhizhong Wang, Lei Zhao, Haibo Chen, Lihong Qiu, Qihang Mo, Sihuan Lin, Wei Xing, Dongming Lu
HIGHLIGHT: In this paper, we tackle these limitations and propose a simple yet effective method for diversified arbitrary style transfer.

772, TITLE: MSG-GAN: Multi-Scale Gradients for Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Karnewar_MSG-GAN_Multi-Scale_Gradients_for_Generative_Adversarial_Networks_CVPR_2020_paper.html
AUTHORS: Animesh Karnewar, Oliver Wang
HIGHLIGHT: In this work, we propose the Multi-Scale Gradient Generative Adversarial Network (MSG-GAN), a simple but effective technique for addressing this by allowing the flow of gradients from the discriminator to the generator at multiple scales.

773, TITLE: Overcoming Multi-Model Forgetting in One-Shot NAS With Diversity Maximization
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Overcoming_Multi-Model_Forgetting_in_One-Shot_NAS_With_Diversity_Maximization_CVPR_2020_paper.html
AUTHORS: Miao Zhang, Huiqi Li, Shirui Pan, Xiaojun Chang, Steven Su
HIGHLIGHT: In this paper, we formulate the supernet training in the One-Shot NAS as a constrained optimization problem of continual learning that the learning of current architecture should not degrade the performance of previous architectures during the supernet training.

774, TITLE: Select to Better Learn: Fast and Accurate Deep Learning Using Data Selection From Nonlinear Manifolds
http://openaccess.thecvf.com/content_CVPR_2020/html/Joneidi_Select_to_Better_Learn_Fast_and_Accurate_Deep_Learning_Using_CVPR_2020_paper.html
AUTHORS: Mohsen Joneidi, Saeed Vahidian, Ashkan Esmaeili, Weijia Wang, Nazanin Rahnavard, Bill Lin, Mubarak Shah
HIGHLIGHT: A simple and efficient selection algorithm with a linear complexity order, referred to as spectrum pursuit (SP), is proposed that pursuits spectral components of the dataset using available sample points.

775, TITLE: Neural Point Cloud Rendering via Multi-Plane Projection
http://openaccess.thecvf.com/content_CVPR_2020/html/Dai_Neural_Point_Cloud_Rendering_via_Multi-Plane_Projection_CVPR_2020_paper.html
AUTHORS: Peng Dai, Yinda Zhang, Zhuwen Li, Shuaicheng Liu, Bing Zeng
HIGHLIGHT: We present a new deep point cloud rendering pipeline through multi-plane projections.

776, TITLE: Wish You Were Here: Context-Aware Human Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Gafni_Wish_You_Were_Here_Context-Aware_Human_Generation_CVPR_2020_paper.html
AUTHORS: Oran Gafni, Lior Wolf
HIGHLIGHT: We present a novel method for inserting objects, specifically humans, into existing images, such that they blend in a photorealistic manner, while respecting the semantic context of the scene.

777, TITLE: Towards Photo-Realistic Virtual Try-On by Adaptively Generating-Preserving Image Content
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Towards_Photo-Realistic_Virtual_Try-On_by_Adaptively_Generating-Preserving_Image_Content_CVPR_2020_paper.html
AUTHORS: Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wangmeng Zuo, Ping Luo
HIGHLIGHT: To address this issue, we propose a novel visual try-on network, namely Adaptive Content Generating and Preserving Network (ACGPN).

778, TITLE: Breaking the Cycle - Colleagues Are All You Need
http://openaccess.thecvf.com/content_CVPR_2020/html/Nizan_Breaking_the_Cycle_-_Colleagues_Are_All_You_Need_CVPR_2020_paper.html
AUTHORS: Ori Nizan, Ayellet Tal
HIGHLIGHT: This paper proposes a novel approach to performing image-to-image translation between unpaired domains.

779, TITLE: Local Class-Specific and Global Image-Level Generative Adversarial Networks for Semantic-Guided Scene Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Local_Class-Specific_and_Global_Image-Level_Generative_Adversarial_Networks_for_Semantic-Guided_CVPR_2020_paper.html
AUTHORS: Hao Tang, Dan Xu, Yan Yan, Philip H.S. Torr, Nicu Sebe
HIGHLIGHT: In this paper, we address the task of semantic-guided scene generation.

780, TITLE: ManiGAN: Text-Guided Image Manipulation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_ManiGAN_Text-Guided_Image_Manipulation_CVPR_2020_paper.html
AUTHORS: Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, Philip H.S. Torr
HIGHLIGHT: The goal of our paper is to semantically edit parts of an image matching a given text that describes desired attributes (e.g., texture, colour, and background), while preserving other contents that are irrelevant to the text.

781, TITLE: Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions
http://openaccess.thecvf.com/content_CVPR_2020/html/Durall_Watch_Your_Up-Convolution_CNN_Based_Generative_Deep_Neural_Networks_Are_CVPR_2020_paper.html
AUTHORS: Ricard Durall, Margret Keuper, Janis Keuper
HIGHLIGHT: In this paper, we show that common up-sampling methods, i.e. known as up-convolution or transposed convolution, are causing the inability of such models to reproduce spectral distributions of natural training data correctly.

782, TITLE: Belief Propagation Reloaded: Learning BP-Layers for Labeling Problems
http://openaccess.thecvf.com/content_CVPR_2020/html/Knobelreiter_Belief_Propagation_Reloaded_Learning_BP-Layers_for_Labeling_Problems_CVPR_2020_paper.html
AUTHORS: Patrick Knobelreiter, Christian Sormann, Alexander Shekhovtsov, Friedrich Fraundorfer, Thomas Pock
HIGHLIGHT: In this work we take one of the simplest inference methods, a truncated max-product Belief Propagation, and add what is necessary to make it a proper component of a deep learning model: connect it to learning formulations with losses on marginals and compute the backprop operation.

783, TITLE: Barycenters of Natural Images Constrained Wasserstein Barycenters for Image Morphing
http://openaccess.thecvf.com/content_CVPR_2020/html/Simon_Barycenters_of_Natural_Images__Constrained_Wasserstein_Barycenters_for_Image_CVPR_2020_paper.html
AUTHORS: Dror Simon, Aviad Aberdam
HIGHLIGHT: In this work, we propose a novel approach for image morphing that possesses all three desired properties.

784, TITLE: Guided Variational Autoencoder for Disentanglement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Ding_Guided_Variational_Autoencoder_for_Disentanglement_Learning_CVPR_2020_paper.html
AUTHORS: Zheng Ding, Yifan Xu, Weijian Xu, Gaurav Parmar, Yang Yang, Max Welling, Zhuowen Tu
HIGHLIGHT: We propose an algorithm, guided variational autoencoder (Guided-VAE), that is able to learn a controllable generative model by performing latent representation disentanglement learning.

785, TITLE:	Cross-Spectral Face Hallucination via Disentangling Independent Factors
http://openaccess.thecvf.com/content_CVPR_2020/html/Duan_Cross-Spectral_Face_Hallucination_via_Disentangling_Independent_Factors_CVPR_2020_paper.html
AUTHORS:	Boyan Duan, Chaoyou Fu, Yi Li, Xingguang Song, Ran He
HIGHLIGHT:	Rather than building a monolithic but complex structure, this paper proposes a Pose Aligned Cross-spectral Hallucination (PACH) approach to disentangle the independent factors and deal with them in individual stages.

786, TITLE:	Learned Image Compression With Discretized Gaussian Mixture Likelihoods and Attention Modules
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Learned_Image_Compression_With_Discretized_Gaussian_Mixture_Likelihoods_and_Attention_CVPR_2020_paper.html
AUTHORS:	Zhengxue Cheng, Heming Sun, Masaru Takeuchi, Jiro Katto
HIGHLIGHT:	Therefore, in this paper, we propose to use discretized Gaussian Mixture Likelihoods to parameterize the distributions of latent codes, which can achieve a more accurate and flexible entropy model.

787, TITLE:	C-Flow: Conditional Generative Flow Models for Images and 3D Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Pumarola_C-Flow_Conditional_Generative_Flow_Models_for_Images_and_3D_Point_CVPR_2020_paper.html
AUTHORS:	Albert Pumarola, Stefan Popov, Francesc Moreno-Noguer, Vittorio Ferrari
HIGHLIGHT:	In this paper, we introduce C-Flow, a novel conditioning scheme that brings normalizing flows to an entirely new scenario with great possibilities for multimodal data modeling.

788, TITLE:	Cogradient Descent for Bilinear Optimization
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhuo_Cogradient_Descent_for_Bilinear_Optimization_CVPR_2020_paper.html
AUTHORS:	Li'an Zhuo, Baochang Zhang, Linlin Yang, Hanlin Chen, Qixiang Ye, David Doermann, Rongrong Ji, Guodong Guo
HIGHLIGHT:	In this paper, we introduce a Cogradient Descent algorithm (CoGD) to address the bilinear problem, based on a theoretical framework to coordinate the gradient of hidden variables via a projection function.

789, TITLE:	Instance-Aware Image Colorization
http://openaccess.thecvf.com/content_CVPR_2020/html/Su_Instance-Aware_Image_Colorization_CVPR_2020_paper.html
AUTHORS:	Jheng-Wei Su, Hung-Kuo Chu, Jia-Bin Huang
HIGHLIGHT:	In this paper, we propose a method for achieving instance-aware colorization.

790, TITLE:	Joint Training of Variational Auto-Encoder and Latent Energy-Based Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_Joint_Training_of_Variational_Auto-Encoder_and_Latent_Energy-Based_Model_CVPR_2020_paper.html
AUTHORS:	Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT:	This paper proposes a joint training method to learn both the variational auto-encoder (VAE) and the latent energy-based model (EBM).

791, TITLE:	Adaptive Loss-Aware Quantization for Multi-Bit Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Qu_Adaptive_Loss-Aware_Quantization_for_Multi-Bit_Networks_CVPR_2020_paper.html
AUTHORS:	Zhongnan Qu, Zimu Zhou, Yun Cheng, Lothar Thiele
HIGHLIGHT:	We propose Adaptive Loss-aware Quantization (ALQ), a new MBN quantization pipeline that is able to achieve an average bitwidth below one-bit without notable loss in inference accuracy.

792, TITLE:	ScopeFlow: Dynamic Scene Scoping for Optical Flow
http://openaccess.thecvf.com/content_CVPR_2020/html/Bar-Haim_ScopeFlow_Dynamic_Scene_Scoping_for_Optical_Flow_CVPR_2020_paper.html
AUTHORS:	Aviram Bar-Haim, Lior Wolf
HIGHLIGHT:	We propose to modify the common training protocols of optical flow, leading to sizable accuracy improvements without adding to the computational complexity of the training process.

793, TITLE:	Video Super-Resolution With Temporal Group Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Isobe_Video_Super-Resolution_With_Temporal_Group_Attention_CVPR_2020_paper.html
AUTHORS:	Takashi Isobe, Songjiang Li, Xu Jia, Shanxin Yuan, Gregory Slabaugh, Chunjing Xu, Ya-Li Li, Shengjin Wang, Qi Tian

HIGHLIGHT: In this work, we propose a novel method that can effectively incorporate temporal information in a hierarchical way.

794, TITLE: Group Sparsity: The Hinge Between Filter Pruning and Decomposition for Network Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Group_Sparsity_The_Hinge_Between_Filter_Pruning_and_Decomposition_for_CVPR_2020_paper.html
AUTHORS: Yawei Li, Shuhang Gu, Christoph Mayer, Luc Van Gool, Radu Timofte
HIGHLIGHT: In this paper, we analyze two popular network compression techniques, i.e. filter pruning and low-rank decomposition, in a unified sense.

795, TITLE: 3D Photography Using Context-Aware Layered Depth Inpainting
http://openaccess.thecvf.com/content_CVPR_2020/html/Shih_3D_Photography_Using_Context-Aware_Layered_Depth_Inpainting_CVPR_2020_paper.html
AUTHORS: Meng-Li Shih, Shih-Yang Su, Johannes Kopf, Jia-Bin Huang
HIGHLIGHT: We propose a method for converting a single RGB-D input image into a 3D photo, i.e., a multi-layer representation for novel view synthesis that contains hallucinated color and depth structures in regions occluded in the original view.

796, TITLE: MixNMatch: Multifactor Disentanglement and Encoding for Conditional Image Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_MixNMatch_Multifactor_Disentanglement_and_Encoding_for_Conditional_Image_Generation_CVPR_2020_paper.html
AUTHORS: Yuheng Li, Krishna Kumar Singh, Utkarsh Ojha, Yong Jae Lee
HIGHLIGHT: We present MixNMatch, a conditional generative model that learns to disentangle and encode background, object pose, shape, and texture from real images with minimal supervision, for mix-and-match image generation.

797, TITLE: Low-Rank Compression of Neural Nets: Learning the Rank of Each Layer
http://openaccess.thecvf.com/content_CVPR_2020/html/Idelbayev_Low-Rank_Compression_of_Neural_Nets_Learning_the_Rank_of_Each_CVPR_2020_paper.html
AUTHORS: Yerlan Idelbayev, Miguel A. Carreira-Perpinan
HIGHLIGHT: We show that, with a suitable formulation, this problem is amenable to a mixed discrete-continuous optimization jointly over the ranks and over the matrix elements, and give a corresponding algorithm.

798, TITLE: Global Texture Enhancement for Fake Face Detection in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Global_Texture_Enhancement_for_Fake_Face_Detection_in_the_Wild_CVPR_2020_paper.html
AUTHORS: Zhengzhe Liu, Xiaojuan Qi, Philip H.S. Torr
HIGHLIGHT: In this paper, we conduct an empirical study on fake/real faces, and have two important observations: firstly, the texture of fake faces is substantially different from real ones; secondly, global texture statistics are more robust to image editing and transferable to fake faces from different GANs and datasets.

799, TITLE: Panoptic-Based Image Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Dundar_Panoptic-Based_Image_Synthesis_CVPR_2020_paper.html
AUTHORS: Aysegul Dundar, Karan Sapra, Guilin Liu, Andrew Tao, Bryan Catanzaro
HIGHLIGHT: We propose a panoptic aware image synthesis network to generate high fidelity and photorealistic images conditioned on panoptic maps which unify semantic and instance information.

800, TITLE: Lighthouse: Predicting Lighting Volumes for Spatially-Coherent Illumination
http://openaccess.thecvf.com/content_CVPR_2020/html/Srinivasan_Lighthouse_Predicting_Lighting_Volumes_for_Spatially-Coherent_Illumination_CVPR_2020_paper.html
AUTHORS: Pratul P. Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T. Barron, Richard Tucker, Noah Snavely
HIGHLIGHT: We present a deep learning solution for estimating the incident illumination at any 3D location within a scene from an input narrow-baseline stereo image pair.

801, TITLE: Learning to Cartoonize Using White-Box Cartoon Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Learning_to_Cartoonize_Using_White-Box_Cartoon_Representations_CVPR_2020_paper.html
AUTHORS: Xinrui Wang, Jinze Yu
HIGHLIGHT: This paper presents an approach for image cartoonization.

802, TITLE: End-to-End Learnable Geometric Vision by Backpropagating PnP Optimization
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_End-to-End_Learnable_Geometric_Vision_by_Backpropagating_PnP_Optimization_CVPR_2020_paper.html

AUTHORS: Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, Tat-Jun Chin
HIGHLIGHT: Towards this aim, we present BPnP, a novel network module that backpropagates gradients through a Perspective-n-Points (PnP) solver to guide parameter updates of a neural network.

803, TITLE: Analyzing and Improving the Image Quality of StyleGAN
http://openaccess.thecvf.com/content_CVPR_2020/html/Karras_Analyzing_and_Improving_the_Image_Quality_of_StyleGAN_CVPR_2020_paper.html
AUTHORS: Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, Timo Aila
HIGHLIGHT: We expose and analyze several of its characteristic artifacts, and propose changes in both model architecture and training methods to address them.

804, TITLE: Fashion Editing With Adversarial Parsing Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Fashion_Editing_With_Adversarial_Parsing_Learning_CVPR_2020_paper.html
AUTHORS: Haoye Dong, Xiaodan Liang, Yixuan Zhang, Xujie Zhang, Xiaohui Shen, Zhenyu Xie, Bowen Wu, Jian Yin
HIGHLIGHT: In this paper, we propose a novel Fashion Editing Generative Adversarial Network (FE-GAN), which is capable of manipulating fashion images by free-form sketches and sparse color strokes.

805, TITLE: Augment Your Batch: Improving Generalization Through Instance Repetition
http://openaccess.thecvf.com/content_CVPR_2020/html/Hoffer_Augment_Your_Batch_Improving_Generalization_Through_Instance_Repetition_CVPR_2020_paper.html
AUTHORS: Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefler, Daniel Soudry
HIGHLIGHT: We propose to use batch augmentation: replicating instances of samples within the same batch with different data augmentations.

806, TITLE: ARShadowGAN: Shadow Generative Adversarial Network for Augmented Reality in Single Light Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_ARShadowGAN_Shadow_Generative_Adversarial_Network_for_Augmented_Reality_in_Single_CVPR_2020_paper.html
AUTHORS: Daquan Liu, Chengjiang Long, Hongpan Zhang, Hanning Yu, Xinzhi Dong, Chunxia Xiao
HIGHLIGHT: To address this problem, we propose an end-to-end Generative Adversarial Network for shadow generation named ARShadowGAN for augmented reality in single light scenes.

807, TITLE: An End-to-End Edge Aggregation Network for Moving Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Patil_An_End-to-End_Edge_Aggregation_Network_for_Moving_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Prashant W. Patil, Kuldeep M. Biradar, Akshay Dudhane, Subrahmanyam Murala
HIGHLIGHT: In this paper, the inherent correlation learning-based edge extraction mechanism (EEM) and dense residual block (DRB) are proposed for the discriminative foreground representation.

808, TITLE: Learning Video Stabilization Using Optical Flow
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Learning_Video_Stabilization_Using_Optical_Flow_CVPR_2020_paper.html
AUTHORS: Jiyang Yu, Ravi Ramamoorthi
HIGHLIGHT: We propose a novel neural network that infers the per-pixel warp fields for video stabilization from the optical flow fields of the input video.

809, TITLE: Reusing Discriminators for Encoding: Towards Unsupervised Image-to-Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Reusing_Discriminators_for_Encoding_Towards_Unsupervised_Image-to-Image_Translation_CVPR_2020_paper.html
AUTHORS: Runfa Chen, Wenbing Huang, Binghui Huang, Fuchun Sun, Bin Fang
HIGHLIGHT: To tackle this issue, we develop a decoupled training strategy by which the encoder is only trained when maximizing the adversary loss while keeping frozen otherwise.

810, TITLE: Robust Design of Deep Neural Networks Against Adversarial Attacks Based on Lyapunov Theory
http://openaccess.thecvf.com/content_CVPR_2020/html/Rahnama_Robust_Design_of_Deep_Neural_Networks_Against_Adversarial_Attacks_Based_CVPR_2020_paper.html
AUTHORS: Arash Rahnama, Andre T. Nguyen, Edward Raff
HIGHLIGHT: In this work, we take a control theoretic approach to the problem of robustness in DNNs.

811, TITLE: StarGAN v2: Diverse Image Synthesis for Multiple Domains

http://openaccess.thecvf.com/content_CVPR_2020/html/Choi_StarGAN_v2_Diverse_Image_Synthesis_for_Multiple_Domains_CVPR_2020_paper.html
AUTHORS:		Yunjey Choi, Youngjung Uh, Jaejun Yoo, Jung-Woo Ha
HIGHLIGHT:		We propose StarGAN v2, a single framework that tackles both and shows significantly improved results over the baselines.

812, TITLE:		Warping Residual Based Image Stitching for Large Parallax
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Warping_Residual_Based_Image_Stitching_for_Large_Parallax_CVPR_2020_paper.html
AUTHORS:		Kyu-Yul Lee, Jae-Young Sim
HIGHLIGHT:		In this paper, we propose an image stitching algorithm robust to large parallax based on the novel concept of warping residuals.

813, TITLE:		A U-Net Based Discriminator for Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Schonfeld_A_U-Net_Based_Discriminator_for_Generative_Adversarial_Networks_CVPR_2020_paper.html
AUTHORS:		Edgar Schonfeld, Bernt Schiele, Anna Khoreva
HIGHLIGHT:		To target this issue we propose an alternative U-Net based discriminator architecture, borrowing the insights from the segmentation literature.

814, TITLE:		Unpaired Portrait Drawing Generation via Asymmetric Cycle Mapping
http://openaccess.thecvf.com/content_CVPR_2020/html/Yi_Unpaired_Portrait_Drawing_Generation_via_Asymmetric_Cycle_Mapping_CVPR_2020_paper.html
AUTHORS:		Ran Yi, Yong-Jin Liu, Yu-Kun Lai, Paul L. Rosin
HIGHLIGHT:		To address this problem, we propose a novel asymmetric cycle mapping that enforces the reconstruction information to be visible (by a truncation loss) and only embedded in selective facial regions (by a relaxed forward cycle-consistency loss).

815, TITLE:		When to Use Convolutional Neural Networks for Inverse Problems
http://openaccess.thecvf.com/content_CVPR_2020/html/Chodosh_When_to_Use_Convolutional_Neural_Networks_for_Inverse_Problems_CVPR_2020_paper.html
AUTHORS:		Nathaniel Chodosh, Simon Lucey
HIGHLIGHT:		In this work we argue that for some types of inverse problems the CNN approximation breaks down leading to poor performance.

816, TITLE:		LUVLi Face Alignment: Estimating Landmarks' Location, Uncertainty, and Visibility Likelihood
http://openaccess.thecvf.com/content_CVPR_2020/html/Kumar_LUVLi_Face_Alignment_Estimating_Landmarks_Location_Uncertainty_and_Visibility_Likelihood_CVPR_2020_paper.html
AUTHORS:		Abhinav Kumar, Tim K. Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherian, Toshiaki Koike-Akino, Xiaoming Liu, Chen Feng
HIGHLIGHT:		In this paper, we present a novel framework for jointly predicting landmark locations, associated uncertainties of these predicted locations, and landmark visibilities.

817, TITLE:		Affinity Graph Supervision for Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Affinity_Graph_Supervision_for_Visual_Recognition_CVPR_2020_paper.html
AUTHORS:		Chu Wang, Babak Samari, Vladimir G. Kim, Siddhartha Chaudhuri, Kaleem Siddiqi
HIGHLIGHT:		Here we propose a principled method to directly supervise the learning of weights in affinity graphs, to exploit meaningful connections between entities in the data source.

818, TITLE:		Unsupervised Magnification of Posture Deviations Across Subjects
http://openaccess.thecvf.com/content_CVPR_2020/html/Dorkenwald_Unsupervised_Magnification_of_Posture_Deviations_Across_Subjects_CVPR_2020_paper.html
AUTHORS:		Michael Dorkenwald, Uta Buchler, Bjorn Ommer
HIGHLIGHT:		We present an approach to unsupervised magnification of posture differences across individuals despite large deviations in appearance.

819, TITLE:		Accurate Estimation of Body Height From a Single Depth Image via a Four-Stage Developing Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Yin_Accurate_Estimation_of_Body_Height_From_a_Single_Depth_Image_CVPR_2020_paper.html
AUTHORS:		Fukun Yin, Shizhe Zhou

HIGHLIGHT: In this paper we address the problem of accurately estimating the height of a person with arbitrary postures from a single depth image.

820, TITLE: Fast Soft Color Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Akimoto_Fast_Soft_Color_Segmentation_CVPR_2020_paper.html
AUTHORS: Naofumi Akimoto, Huachun Zhu, Yanghua Jin, Yoshimitsu Aoki
HIGHLIGHT: To address this issue, we propose a neural network based method for this task that decomposes a given image into multiple layers in a single forward pass.

821, TITLE: Global Optimality for Point Set Registration Using Semidefinite Programming
http://openaccess.thecvf.com/content_CVPR_2020/html/Iglesias_Global_Optimality_for_Point_Set_Registration_Using_Semidefinite_Programming_CVPR_2020_paper.html
AUTHORS: Jose Pedro Iglesias, Carl Olsson, Fredrik Kahl
HIGHLIGHT: In this paper we present a study of global optimality conditions for Point Set Registration (PSR) with missing data.

822, TITLE: Image2StyleGAN++: How to Edit the Embedded Images?
http://openaccess.thecvf.com/content_CVPR_2020/html/Abdal_Image2StyleGAN_How_to_Edit_the_Embedded_Images_CVPR_2020_paper.html
AUTHORS: Rameen Abdal, Yipeng Qin, Peter Wonka
HIGHLIGHT: We propose Image2StyleGAN++, a flexible image editing framework with many applications.

823, TITLE: SQE: a Self Quality Evaluation Metric for Parameters Optimization in Multi-Object Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_SQE_a_Self_Quality_Evaluation_Metric_for_Parameters_Optimization_in_CVPR_2020_paper.html
AUTHORS: Yanru Huang, Feiyu Zhu, Zheni Zeng, Xi Qiu, Yuan Shen, Jianan Wu
HIGHLIGHT: We present a novel self quality evaluation metric SQE for parameters optimization in the challenging yet critical multi-object tracking task.

824, TITLE: EventSR: From Asynchronous Events to Image Reconstruction, Restoration, and Super-Resolution via End-to-End Adversarial Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_EventSR_From_Asynchronous_Events_to_Image_Reconstruction_Restoration_and_Super-Resolution_CVPR_2020_paper.html
AUTHORS: Lin Wang, Tae-Kyun Kim, Kuk-Jin Yoon
HIGHLIGHT: To tackle the challenges, we propose a novel end-to-end pipeline that reconstructs LR images from event streams, enhances the image qualities and upsamples the enhanced images, called EventSR.

825, TITLE: Hierarchical Pyramid Diverse Attention Networks for Face Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Hierarchical_Pyramid_Diverse_Attention_Networks_for_Face_Recognition_CVPR_2020_paper.html
AUTHORS: Qiangchang Wang, Tianyi Wu, He Zheng, Guodong Guo
HIGHLIGHT: In this work, we propose a hierarchical pyramid diverse attention (HPDA) network.

826, TITLE: RGBD-Dog: Predicting Canine Pose from RGBD Sensors
http://openaccess.thecvf.com/content_CVPR_2020/html/Kearney_RGBD-Dog_Predicting_Canine_Pose_from_RGBD_Sensors_CVPR_2020_paper.html
AUTHORS: Sinead Kearney, Wenbin Li, Martin Parsons, Kwang In Kim, Darren Cosker
HIGHLIGHT: In our work, we focus on the problem of 3D canine pose estimation from RGBD images, recording a diverse range of dog breeds with several Microsoft Kinect v2s, simultaneously obtaining the 3D ground truth skeleton via a motion capture system.

827, TITLE: Multi-Scale Progressive Fusion Network for Single Image Deraining
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_Multi-Scale_Progressive_Fusion_Network_for_Single_Image_Deraining_CVPR_2020_paper.html
AUTHORS: Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, Junjun Jiang
HIGHLIGHT: In this work, we explore the multi-scale collaborative representation for rain streaks from the perspective of input image scales and hierarchical deep features in a unified framework, termed multi-scale progressive fusion network (MSPFN) for single image rain streak removal.

828, TITLE: Learning a Neural 3D Texture Space From 2D Exemplars

http://openaccess.thecvf.com/content_CVPR_2020/html/Henzler_Learning_a_Neural_3D_Texture_Space_From_2D_Exemplars_CVPR_2020_paper.html
AUTHORS: Philipp Henzler, Niloy J. Mitra, Tobias Ritschel
HIGHLIGHT: We suggest a generative model of 2D and 3D natural textures with diversity, visual fidelity and at high computational efficiency.

829, TITLE: BachGAN: High-Resolution Image Synthesis From Salient Object Layout
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_BachGAN_High-Resolution_Image_Synthesis_From_Salient_Object_Layout_CVPR_2020_paper.html
AUTHORS: Yandong Li, Yu Cheng, Zhe Gan, Licheng Yu, Liqiang Wang, Jingjing Liu
HIGHLIGHT: We propose a new task towards more practical applications for image generation - high-quality image synthesis from salient object layout.

830, TITLE: Rethinking Data Augmentation for Image Super-resolution: A Comprehensive Analysis and a New Strategy
http://openaccess.thecvf.com/content_CVPR_2020/html/Yoo_Rethinking_Data_Augmentation_for_Image_Super-resolution_A_Comprehensive_Analysis_and_CVPR_2020_paper.html
AUTHORS: Jaejun Yoo, Namhyuk Ahn, Kyung-Ah Sohn
HIGHLIGHT: In this paper, we provide a comprehensive analysis of the existing augmentation methods applied to the super-resolution task.

831, TITLE: On Positive-Unlabeled Classification in GAN
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_On_Positive-Unlabeled_Classification_in_GAN_CVPR_2020_paper.html
AUTHORS: Tianyu Guo, Chang Xu, Jiajun Huang, Yunhe Wang, Boxin Shi, Chao Xu, Dacheng Tao
HIGHLIGHT: This paper defines a positive and unlabeled classification problem for standard GANs, which then leads to a novel technique to stabilize the training of the discriminator in GANs.

832, TITLE: DoveNet: Deep Image Harmonization via Domain Verification
http://openaccess.thecvf.com/content_CVPR_2020/html/Cong_DoveNet_Deep_Image_Harmonization_via_Domain_Verification_CVPR_2020_paper.html
AUTHORS: Wenyan Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, Liqing Zhang
HIGHLIGHT: In this work, we contribute an image harmonization dataset iHarmony4 by generating synthesized composite images based on COCO (resp., Adobe5k, Flickr, day2night) dataset, leading to our HCOCO (resp., HAdobe5k, HFlickr, Hday2night) sub-dataset.

833, TITLE: Noise Robust Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Kaneko_Noise_Robust_Generative_Adversarial_Networks_CVPR_2020_paper.html
AUTHORS: Takuhiro Kaneko, Tatsuya Harada
HIGHLIGHT: As an alternative, we propose a novel family of GANs called noise robust GANs (NR-GANs), which can learn a clean image generator even when training images are noisy.

834, TITLE: Normalizing Flows With Multi-Scale Autoregressive Priors
http://openaccess.thecvf.com/content_CVPR_2020/html/Bhattacharyya_Normalizing_Flows_With_Multi-Scale_Autoregressive_Priors_CVPR_2020_paper.html
AUTHORS: Apratim Bhattacharyya, Shweta Mahajan, Mario Fritz, Bernt Schiele, Stefan Roth
HIGHLIGHT: In this work, we improve the representational power of flow-based models by introducing channel-wise dependencies in their latent space through multi-scale autoregressive priors (mAR).

835, TITLE: Robust Reference-Based Super-Resolution With Similarity-Aware Deformable Convolution
http://openaccess.thecvf.com/content_CVPR_2020/html/Shim_Robust_Reference-Based_Super-Resolution_With_Similarity-Aware_Deformable_Convolution_CVPR_2020_paper.html
AUTHORS: Gyumin Shim, Jinsun Park, In So Kweon
HIGHLIGHT: In this paper, we propose a novel and efficient reference feature extraction module referred to as the Similarity Search and Extraction Network (SSEN) for reference-based super-resolution (RefSR) tasks.

836, TITLE: Painting Many Pasts: Synthesizing Time Lapse Videos of Paintings
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Painting_Many_Pasts_Synthesizing_Time_Lapse_Videos_of_Paintings_CVPR_2020_paper.html
AUTHORS: Amy Zhao, Guha Balakrishnan, Kathleen M. Lewis, Fredo Durand, John V. Guttag, Adrian V. Dalca
HIGHLIGHT: We present a probabilistic model that, given a single image of a completed painting, recurrently synthesizes steps of the painting process.

837, TITLE: GeoDA: A Geometric Framework for Black-Box Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2020/html/Rahmati_GeoDA_A_Geometric_Framework_for_Black-Box_Adversarial_Attacks_CVPR_2020_paper.html
AUTHORS: Ali Rahmati, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard, Huaiyu Dai
HIGHLIGHT: We propose a geometric framework to generate adversarial examples in one of the most challenging black-box settings where the adversary can only generate a small number of queries, each of them returning the top-1 label of the classifier.

838, TITLE: GAMIN: Generative Adversarial Multiple Imputation Network for Highly Missing Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Yoon_GAMIN_Generative_Adversarial_Multiple_Imputation_Network_for_Highly_Missing_Data_CVPR_2020_paper.html
AUTHORS: Seongwook Yoon, Sanghoon Sull
HIGHLIGHT: We propose a novel imputation method for highly missing data.

839, TITLE: An Internal Covariate Shift Bounding Algorithm for Deep Neural Networks by Unitizing Layers' Outputs
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_An_Internal_Covariate_Shift_Bounding_Algorithm_for_Deep_Neural_Networks_CVPR_2020_paper.html
AUTHORS: You Huang, Yuanlong Yu
HIGHLIGHT: Thus this paper proposes a measure for ICS by using the Earth Mover (EM) distance and then derives the upper and lower bounds for the measure to provide a theoretical analysis of BN.

840, TITLE: A Unified Optimization Framework for Low-Rank Inducing Penalties
http://openaccess.thecvf.com/content_CVPR_2020/html/Ornhag_A_Unified_Optimization_Framework_for_Low-Rank_Inducing_Penalties_CVPR_2020_paper.html
AUTHORS: Marcus Valtonen Ornhag, Carl Olsson
HIGHLIGHT: In this paper we study the convex envelopes of a new class of functions.

841, TITLE: Single-Side Domain Generalization for Face Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2020/html/Jia_Single-Side_Domain_Generalization_for_Face_Anti-Spoofing_CVPR_2020_paper.html
AUTHORS: Yunpei Jia, Jie Zhang, Shiguang Shan, Xilin Chen
HIGHLIGHT: In this work, we propose an end-to-end single-side domain generalization framework (SSDG) to improve the generalization ability of face anti-spoofing.

842, TITLE: The Knowledge Within: Methods for Data-Free Model Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Haroush_The_Knowledge_Within_Methods_for_Data-Free_Model_Compression_CVPR_2020_paper.html
AUTHORS: Matan Haroush, Itay Hubara, Elad Hoffer, Daniel Soudry
HIGHLIGHT: Contributions: We present three methods for generating synthetic samples from trained models. Then, we demonstrate how these samples can be used to calibrate and fine-tune quantized models without using any real data in the process.

843, TITLE: Scale-Space Flow for End-to-End Optimized Video Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Agustsson_Scale-Space_Flow_for_End-to-End_Optimized_Video_Compression_CVPR_2020_paper.html
AUTHORS: Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, George Toderici
HIGHLIGHT: In this paper, we show that a generalized warping operator that better handles common failure cases, e.g. disocclusions and fast motion, can provide competitive compression results with a greatly simplified model and training procedure.

844, TITLE: Dynamic Neural Relational Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Graber_Dynamic_Neural_Relational_Inference_CVPR_2020_paper.html
AUTHORS: Colin Graber, Alexander G. Schwing
HIGHLIGHT: In response to this, we develop Dynamic Neural Relational Inference (dNRI), which incorporates insights from sequential latent variable models to predict separate relation graphs for every time-step.

845, TITLE: Real-Time Panoptic Segmentation From Dense Detections
http://openaccess.thecvf.com/content_CVPR_2020/html/Hou_Real-Time_Panoptic_Segmentation_From_Dense_Detections_CVPR_2020_paper.html
AUTHORS: Rui Hou, Jie Li, Arjun Bhargava, Allan Raventos, Vitor Guizilini, Chao Fang, Jerome Lynch, Adrien Gaidon
HIGHLIGHT: In this paper, we propose a new single-shot panoptic segmentation network that leverages dense detections and a global self-attention mechanism to operate in real-time with performance approaching the state of the art.

846, TITLE: Deep Snake for Real-Time Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Peng_Deep_Snake_for_Real-Time_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Sida Peng, Wen Jiang, Huaijin Pi, Xiuli Li, Hujun Bao, Xiaowei Zhou
HIGHLIGHT: This paper introduces a novel contour-based approach named deep snake for real-time instance segmentation.

847, TITLE: AdaCoSeg: Adaptive Shape Co-Segmentation With Group Consistency Loss
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_AdaCoSeg_Adaptive_Shape_Co-Segmentation_With_Group_Consistency_Loss_CVPR_2020_paper.html
AUTHORS: Chenyang Zhu, Kai Xu, Siddhartha Chaudhuri, Li Yi, Leonidas J. Guibas, Hao Zhang
HIGHLIGHT: We introduce AdaCoSeg, a deep neural network architecture for adaptive co-segmentation of a set of 3D shapes represented as point clouds.

848, TITLE: Learning Dynamic Routing for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Learning_Dynamic_Routing_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Yanwei Li, Lin Song, Yukang Chen, Zeming Li, Xiangyu Zhang, Xingang Wang, Jian Sun
HIGHLIGHT: This paper studies a conceptually new method to alleviate the scale variance in semantic representation, named dynamic routing.

849, TITLE: Boosting Semantic Human Matting With Coarse Annotations
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Boosting_Semantic_Human_Matting_With_Coarse_Annotations_CVPR_2020_paper.html
AUTHORS: Jinlin Liu, Yuan Yao, Wendi Hou, Miaomiao Cui, Xuansong Xie, Changshui Zhang, Xian-Sheng Hua
HIGHLIGHT: In this paper, we propose to leverage coarse annotated data coupled with fine annotated data to boost end-to-end semantic human matting without trimaps as extra input.

850, TITLE: BlendMask: Top-Down Meets Bottom-Up for Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_BlendMask_Top-Down_Meets_Bottom-Up_for_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, Youliang Yan
HIGHLIGHT: Our main contribution is a blender module which draws inspiration from both top-down and bottom-up instance segmentation approaches.

851, TITLE: UC-Net: Uncertainty Inspired RGB-D Saliency Detection via Conditional Variational Autoencoders
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_UC-Net_Uncertainty_Inspired_RGB-D_Saliency_Detection_via_Conditional_Variational_Autoencoders_CVPR_2020_paper.html
AUTHORS: Jing Zhang, Deng-Ping Fan, Yuchao Dai, Saeed Anwar, Fatemeh Sadat Saleh, Tong Zhang, Nick Barnes
HIGHLIGHT: In this paper, we propose the first framework (UCNet) to employ uncertainty for RGB-D saliency detection by learning from the data labeling process.

852, TITLE: Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence
http://openaccess.thecvf.com/content_CVPR_2020/html/Donati_Deep_Geometric_Functional_Maps_Robust_Feature_Learning_for_Shape_Correspondence_CVPR_2020_paper.html
AUTHORS: Nicolas Donati, Abhishek Sharma, Maks Ovsjanikov
HIGHLIGHT: We present a novel learning-based approach for computing correspondences between non-rigid 3D shapes.

853, TITLE: Deep Polarization Cues for Transparent Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kalra_Deep_Polarization_Cues_for_Transparent_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Agastya Kalra, Vage Taamazyan, Supreeth Krishna Rao, Kartik Venkataraman, Ramesh Raskar, Achuta Kadambi
HIGHLIGHT: This paper reframes the problem of transparent object segmentation into the realm of light polarization, i.e., the rotation of light waves.

854, TITLE: DualConvMesh-Net: Joint Geodesic and Euclidean Convolutions on 3D Meshes
http://openaccess.thecvf.com/content_CVPR_2020/html/Schult_DualConvMesh-Net_Joint_Geodesic_and_Euclidean_Convolutions_on_3D_Meshes_CVPR_2020_paper.html
AUTHORS: Jonas Schult, Francis Engelmann, Theodora Kontogianni, Bastian Leibe
HIGHLIGHT: We propose DualConvMesh-Nets (DCM-Net) a family of deep hierarchical convolutional networks over 3D geometric data that combines two types of convolutions.

855, TITLE: F-BRS: Rethinking Backpropagating Refinement for Interactive Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Sofiiuk_F-BRS_Rethinking_Backpropagating_Refinement_for_Interactive_Segmentation_CVPR_2020_paper.html
AUTHORS: Konstantin Sofiiuk, Ilia Petrov, Olga Barinova, Anton Konushin
HIGHLIGHT: We propose f-BRS (feature backpropagating refinement scheme) that solves an optimization problem with respect to auxiliary variables instead of the network inputs, and requires running forward and backward passes just for a small part of a network.

856, TITLE: Approximating shapes in images with low-complexity polygons
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Approximating_shapes_in_images_with_low-complexity_polygons_CVPR_2020_paper.html
AUTHORS: Muxingzi Li, Florent Lafarge, Renaud Marlet
HIGHLIGHT: We present an algorithm for extracting and vectorizing objects in images with polygons.

857, TITLE: Towards Visually Explaining Variational Autoencoders
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Towards_Visually_Explaining_Variational_Autoencoders_CVPR_2020_paper.html
AUTHORS: Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyan Wu, Bir Bhanu, Richard J. Radke, Octavia Camps
HIGHLIGHT: In this work, we take a step towards bridging this crucial gap, proposing the first technique to visually explain VAEs by means of gradient-based attention.

858, TITLE: Towards Global Explanations of Convolutional Neural Networks With Concept Attribution
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Towards_Global_Explanations_of_Convolutional_Neural_Networks_With_Concept_Attribution_CVPR_2020_paper.html
AUTHORS: Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R. Lyu, Yu-Wing Tai
HIGHLIGHT: To overcome such drawbacks, we propose a novel two-stage framework, Attacking for Interpretability (AfI), which explains model decisions in terms of the importance of user-defined concepts.

859, TITLE: Interpretable and Accurate Fine-grained Recognition via Region Grouping
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Interpretable_and_Accurate_Fine-grained_Recognition_via_Region_Grouping_CVPR_2020_paper.html
AUTHORS: Zixuan Huang, Yin Li
HIGHLIGHT: We present an interpretable deep model for fine-grained visual recognition.

860, TITLE: SAM: The Sensitivity of Attribution Methods to Hyperparameters
http://openaccess.thecvf.com/content_CVPR_2020/html/Bansal_SAM_The_Sensitivity_of_Attribution_Methods_to_Hyperparameters_CVPR_2020_paper.html
AUTHORS: Naman Bansal, Chirag Agarwal, Anh Nguyen
HIGHLIGHT: In this paper, we provide a thorough empirical study on the sensitivity of existing attribution methods.

861, TITLE: High-Frequency Component Helps Explain the Generalization of Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_High-Frequency_Component_Helps_Explain_the_Generalization_of_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Haohan Wang, Xindi Wu, Zeyi Huang, Eric P. Xing
HIGHLIGHT: We investigate the relationship between the frequency spectrum of image data and the generalization behavior of convolutional neural networks (CNN).

862, TITLE: FALCON: A Fourier Transform Based Approach for Fast and Secure Convolutional Neural Network Predictions
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_FALCON_A_Fourier_Transform_Based_Approach_for_Fast_and_Secure_CVPR_2020_paper.html
AUTHORS: Shaohua Li, Kaiping Xue, Bin Zhu, Chenkai Ding, Xindi Gao, David Wei, Tao Wan
HIGHLIGHT: In this paper, we focus on the scenario where clients want to classify private images with a convolutional neural network model hosted in the server, while both parties keep their data private.

863, TITLE: Dreaming to Distill: Data-Free Knowledge Transfer via DeepInversion
http://openaccess.thecvf.com/content_CVPR_2020/html/Yin_Dreaming_to_Distill_Data-Free_Knowledge_Transfer_via_DeepInversion_CVPR_2020_paper.html

AUTHORS: Hongxu Yin, Pavlo Molchanov, Jose M. Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K. Jha, Jan Kautz
HIGHLIGHT: We introduce DeepInversion, a new method for synthesizing images from the image distribution used to train a deep neural network.

864, TITLE: Unsupervised Domain Adaptation via Structurally Regularized Deep Clustering
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Unsupervised_Domain_Adaptation_via_Structurally_Regularized_Deep_Clustering_CVPR_2020_paper.html
AUTHORS: Hui Tang, Ke Chen, Kui Jia
HIGHLIGHT: To alleviate this risk, we are motivated by the assumption of structural domain similarity, and propose to directly uncover the intrinsic target discrimination via discriminative clustering of target data.

865, TITLE: HyperSTAR: Task-Aware Hyperparameters for Deep Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Mittal_HyperSTAR_Task-Aware_Hyperparameters_for_Deep_Networks_CVPR_2020_paper.html
AUTHORS: Gaurav Mittal, Chang Liu, Nikolaos Karianakis, Victor Fragoso, Mei Chen, Yun Fu
HIGHLIGHT: To reduce HPO time, we present HyperSTAR (System for Task Aware Hyperparameter Recommendation), a task-aware method to warm-start HPO for deep neural networks.

866, TITLE: ActBERT: Learning Global-Local Video-Text Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_ActBERT_Learning_Global-Local_Video-Text_Representations_CVPR_2020_paper.html
AUTHORS: Linchao Zhu, Yi Yang
HIGHLIGHT: In this paper, we introduce ActBERT for self-supervised learning of joint video-text representations from unlabeled data.

867, TITLE: State-Relabeling Adversarial Active Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_State-Relabeling_Adversarial_Active_Learning_CVPR_2020_paper.html
AUTHORS: Beichen Zhang, Liang Li, Shijie Yang, Shuhui Wang, Zheng-Jun Zha, Qingming Huang
HIGHLIGHT: In this paper, we propose a state relabeling adversarial active learning model (SRAAL), that leverages both the annotation and the labeled/unlabeled state information for deriving the most informative unlabeled samples.

868, TITLE: Erasing Integrated Learning: A Simple Yet Effective Approach for Weakly Supervised Object Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Mai_Erasing_Integrated_Learning_A_Simple_Yet_Effective_Approach_for_Weakly_CVPR_2020_paper.html
AUTHORS: Jinjie Mai, Meng Yang, Wenfeng Luo
HIGHLIGHT: To remedy this, we propose a simple yet powerful approach by introducing a novel adversarial erasing technique, erasing integrated learning (EIL).

869, TITLE: A Shared Multi-Attention Framework for Multi-Label Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Huynh_A_Shared_Multi-Attention_Framework_for_Multi-Label_Zero-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Dat Huynh, Ehsan Elhamifar
HIGHLIGHT: In this work, we develop a shared multi-attention model for multi-label zero-shot learning.

870, TITLE: Self-Supervised Learning of Interpretable Keypoints From Unlabelled Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Jakab_Self-Supervised_Learning_of_Interpretable_Keypoints_From_Unlabelled_Videos_CVPR_2020_paper.html
AUTHORS: Tomas Jakab, Ankush Gupta, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: We propose a new method for recognizing the pose of objects from a single image that for learning uses only unlabelled videos and a weak empirical prior on the object poses.

871, TITLE: Few-Shot Open-Set Recognition Using Meta-Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Few-Shot_Open-Set_Recognition_Using_Meta-Learning_CVPR_2020_paper.html
AUTHORS: Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, Nuno Vasconcelos
HIGHLIGHT: This combines the random selection of a set of novel classes per episode, a loss that maximizes the posterior entropy for examples of those classes, and a new metric learning formulation based on the Mahalanobis distance.

872, TITLE: Few-Shot Learning via Embedding Adaptation With Set-to-Set Functions

http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_Few-Shot_Learning_via_Embedding_Adaptation_With_Set-to-Set_Functions_CVPR_2020_paper.html
AUTHORS: Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, Fei Sha
HIGHLIGHT: In this paper, we propose a novel approach to adapt the instance embeddings to the target classification task with a set-to-set function, yielding embeddings that are task-specific and are discriminative.

873, TITLE: Temporally Distributed Networks for Fast Video Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Temporally_Distributed_Networks_for_Fast_Video_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Ping Hu, Fabian Caba, Oliver Wang, Zhe Lin, Stan Sclaroff, Federico Perazzi
HIGHLIGHT: We present TDNet, a temporally distributed network designed for fast and accurate video semantic segmentation.

874, TITLE: Benchmarking the Robustness of Semantic Segmentation Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Kamann_Benchmarking_the_Robustness_of_Semantic_Segmentation_Models_CVPR_2020_paper.html
AUTHORS: Christoph Kamann, Carsten Rother
HIGHLIGHT: While there are recent robustness studies for full-image classification, we are the first to present an exhaustive study for semantic segmentation, based on the state-of-the-art model DeepLabv3+.

875, TITLE: There and Back Again: Revisiting Backpropagation Saliency Methods
http://openaccess.thecvf.com/content_CVPR_2020/html/Rebuffi_There_and_Back_Again_Revisiting_Backpropagation_Saliency_Methods_CVPR_2020_paper.html
AUTHORS: Sylvestre-Alvise Rebuffi, Ruth Fong, Xu Ji, Andrea Vedaldi
HIGHLIGHT: In this work, we conduct a thorough analysis of backpropagation-based saliency methods and propose a single framework under which several such methods can be unified.

876, TITLE: Deep Semantic Clustering by Partition Confidence Maximisation
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Deep_Semantic_Clustering_by_Partition_Confidence_Maximisation_CVPR_2020_paper.html
AUTHORS: Jiabo Huang, Shaogang Gong, Xiatian Zhu
HIGHLIGHT: In this work, we propose to solve this problem by learning the most confident clustering solution from all the possible separations, based on the observation that assigning samples from the same semantic categories into different clusters will reduce both the intra-cluster compactness and inter-cluster diversity, i.e. lower partition confidence.

877, TITLE: StructEdit: Learning Structural Shape Variations
http://openaccess.thecvf.com/content_CVPR_2020/html/Mo_StructEdit_Learning_Structural_Shape_Variations_CVPR_2020_paper.html
AUTHORS: Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy J. Mitra, Leonidas J. Guibas
HIGHLIGHT: Instead, we treat shape differences as primary objects in their own right and propose to encode them in their own latent space.

878, TITLE: Harmonizing Transferability and Discriminability for Adapting Object Detectors
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Harmonizing_Transferability_and_Discriminability_for_Adapting_Object_Detectors_CVPR_2020_paper.html
AUTHORS: Chaoqi Chen, Zebiao Zheng, Xinghao Ding, Yue Huang, Qi Dou
HIGHLIGHT: In this paper, we propose a Hierarchical Transferability Calibration Network (HTCN) that hierarchically (local-region/image/instance) calibrates the transferability of feature representations for harmonizing transferability and discriminability.

879, TITLE: Fast Video Object Segmentation With Temporal Aggregation Network and Dynamic Template Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Fast_Video_Object_Segmentation_With_Temporal_Aggregation_Network_and_Dynamic_CVPR_2020_paper.html
AUTHORS: Xuhua Huang, Jiarui Xu, Yu-Wing Tai, Chi-Keung Tang
HIGHLIGHT: In this paper, we introduce "tracking-by-detection" into VOS which can coherently integrates segmentation into tracking, by proposing a new temporal aggregation network and a novel dynamic time-evolving template matching mechanism to achieve significantly improved performance.

880, TITLE: CascadePSP: Toward Class-Agnostic and Very High-Resolution Segmentation via Global and Local Refinement
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_CascadePSP_Toward_Class-Agnostic_and_Very_High-Resolution_Segmentation_via_Global_and_CVPR_2020_paper.html
AUTHORS: Ho Kei Cheng, Jihoon Chung, Yu-Wing Tai, Chi-Keung Tang

HIGHLIGHT: In this paper, we propose a novel approach to address the high-resolution segmentation problem without using any high-resolution training data.

881, TITLE: Correlating Edge, Pose With Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Correlating_Edge_Pose_With_Parsing_CVPR_2020_paper.html
AUTHORS: Ziwei Zhang, Chi Su, Liang Zheng, Xiaodong Xie
HIGHLIGHT: To capture such correlations, we propose a Correlation Parsing Machine (CorrPM) employing a heterogeneous non-local block to discover the spatial affinity among feature maps from the edge, pose and parsing.

882, TITLE: VecRoad: Point-Based Iterative Graph Exploration for Road Graphs Extraction
http://openaccess.thecvf.com/content_CVPR_2020/html/Tan_VecRoad_Point-Based_Iterative_Graph_Exploration_for_Road_Graphs_Extraction_CVPR_2020_paper.html
AUTHORS: Yong-Qiang Tan, Shang-Hua Gao, Xuan-Yi Li, Ming-Ming Cheng, Bo Ren
HIGHLIGHT: To enhance the road connectivity while maintaining the precise alignment between the graph and real road, we propose a point-based iterative graph exploration scheme with segmentation-cues guidance and flexible steps.

883, TITLE: Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Towards_Fairness_in_Visual_Recognition_Effective_Strategies_for_Bias_Mitigation_CVPR_2020_paper.html
AUTHORS: Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, Olga Russakovsky
HIGHLIGHT: We highlight the shortcomings of popular adversarial training approaches for bias mitigation, propose a simple but similarly effective alternative to the inference-time Reducing Bias Amplification method of Zhao et al., and design a domain-independent training technique that outperforms all other methods.

884, TITLE: Hierarchical Human Parsing With Typed Part-Relation Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Hierarchical_Human_Parsing_With_Typed_Part-Relation_Reasoning_CVPR_2020_paper.html
AUTHORS: Wenguan Wang, Hailong Zhu, Jifeng Dai, Yanwei Pang, Jianbing Shen, Ling Shao
HIGHLIGHT: Focusing on this, we seek to simultaneously exploit the representational capacity of deep graph networks and the hierarchical human structures.

885, TITLE: Compositional Convolutional Neural Networks: A Deep Architecture With Innate Robustness to Partial Occlusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Kortylewski_Compositional_Convolutional_Neural_Networks_A_Deep_Architecture_With_Innate_Robustness_CVPR_2020_paper.html
AUTHORS: Adam Kortylewski, Ju He, Qing Liu, Alan L. Yuille
HIGHLIGHT: Inspired by the success of compositional models at classifying partially occluded objects, we propose to integrate compositional models and DCNNs into a unified deep model with innate robustness to partial occlusion.

886, TITLE: Spatial Pyramid Based Graph Reasoning for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Spatial_Pyramid_Based_Graph_Reasoning_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Xia Li, Yibo Yang, Qijie Zhao, Tiancheng Shen, Zhouchen Lin, Hong Liu
HIGHLIGHT: In this paper, we apply graph convolution into the semantic segmentation task and propose an improved Laplacian.

887, TITLE: Learning Video Object Segmentation From Unlabeled Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Learning_Video_Object_Segmentation_From_Unlabeled_Videos_CVPR_2020_paper.html
AUTHORS: Xiankai Lu, Wenguan Wang, Jianbing Shen, Yu-Wing Tai, David J. Crandall, Steven C. H. Hoi
HIGHLIGHT: We propose a new method for video object segmentation (VOS) that addresses object pattern learning from unlabeled videos, unlike most existing methods which rely heavily on extensive annotated data.

888, TITLE: Part-Aware Context Network for Human Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Part-Aware_Context_Network_for_Human_Parsing_CVPR_2020_paper.html
AUTHORS: Xiaomei Zhang, Yingying Chen, Bingke Zhu, Jinqiao Wang, Ming Tang
HIGHLIGHT: In this work, we propose a Part-aware Context Network (PCNet), a novel and effective algorithm to deal with the challenge.

889, TITLE: SCOUT: Self-Aware Discriminant Counterfactual Explanations
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_SCOUT_Self-Aware_Discriminant_Counterfactual_Explanations_CVPR_2020_paper.html
AUTHORS: Pei Wang, Nuno Vasconcelos
HIGHLIGHT: A new family of discriminant explanations is introduced. These produce heatmaps that attribute high scores to image regions informative of a classifier prediction but not of a counter class.

890, TITLE: Weakly-Supervised Semantic Segmentation via Sub-Category Exploration
http://openaccess.thecvf.com/content_CVPR_2020/html/Chang_Weakly-Supervised_Semantic_Segmentation_via_Sub-Category_Exploration_CVPR_2020_paper.html
AUTHORS: Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, Ming-Hsuan Yang
HIGHLIGHT: To enforce the network to pay attention to other parts of an object, we propose a simple yet effective approach that introduces a self-supervised task by exploiting the sub-category information.

891, TITLE: Continual Learning With Extended Kronecker-Factored Approximate Curvature
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Continual_Learning_With_Extended_Kronecker-Factored_Approximate_Curvature_CVPR_2020_paper.html
AUTHORS: Janghyeon Lee, Hyeong Gwon Hong, Donggyu Joo, Junmo Kim
HIGHLIGHT: We propose a quadratic penalty method for continual learning of neural networks that contain batch normalization (BN) layers.

892, TITLE: Phase Consistent Ecological Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Phase_Consistent_Ecological_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Yanchao Yang, Dong Lao, Ganesh Sundaramoorthi, Stefano Soatto
HIGHLIGHT: We introduce two criteria to regularize the optimization involved in learning a classifier in a domain where no annotated data are available, leveraging annotated data in a different domain, a problem known as unsupervised domain adaptation.

893, TITLE: AD-Cluster: Augmented Discriminative Clustering for Domain Adaptive Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhai_AD-Cluster_Augmented_Discriminative_Clustering_for_Domain_Adaptive_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS: Yunpeng Zhai, Shijian Lu, Qixiang Ye, Xuebo Shan, Jie Chen, Rongrong Ji, Yonghong Tian
HIGHLIGHT: This paper presents a novel augmented discriminative clustering (AD-Cluster) technique that estimates and augments person clusters in target domains and enforces the discrimination ability of re-ID models with the augmented clusters.

894, TITLE: 3D-MPA: Multi-Proposal Aggregation for 3D Semantic Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Engelmann_3D-MPA_Multi-Proposal_Aggregation_for_3D_Semantic_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Francis Engelmann, Martin Bokeloh, Alireza Fathi, Bastian Leibe, Matthias Niessner
HIGHLIGHT: We present 3D-MPA, a method for instance segmentation on 3D point clouds.

895, TITLE: Deep Active Learning for Biased Datasets via Fisher Kernel Self-Supervision
http://openaccess.thecvf.com/content_CVPR_2020/html/Gudovskiy_Deep_Active_Learning_for_Biased_Datasets_via_Fisher_Kernel_Self-Supervision_CVPR_2020_paper.html
AUTHORS: Denis Gudovskiy, Alec Hodgkinson, Takuya Yamaguchi, Sotaro Tsukizawa
HIGHLIGHT: To implement such acquisition function, we propose a low-complexity method for feature density matching using self-supervised Fisher kernel (FK) as well as several novel pseudo-label estimators.

896, TITLE: Adaptive Graph Convolutional Network With Attention Graph Clustering for Co-Saliency Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Adaptive_Graph_Convolutional_Network_With_Attention_Graph_Clustering_for_Co-Saliency_CVPR_2020_paper.html
AUTHORS: Kaihua Zhang, Tengpeng Li, Shiwen Shen, Bo Liu, Jin Chen, Qingshan Liu
HIGHLIGHT: For this task, we present a novel adaptive graph convolutional network with attention graph clustering (GCAGC).

897, TITLE: A2dele: Adaptive and Attentive Depth Distiller for Efficient RGB-D Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Piao_A2dele_Adaptive_and_Attentive_Depth_Distiller_for_Efficient_RGB-D_Salient_CVPR_2020_paper.html
AUTHORS: Yongri Piao, Zhengkun Rong, Miao Zhang, Weisong Ren, Huchuan Lu
HIGHLIGHT: To tackle these two dilemmas, we propose a depth distiller (A2dele) to explore the way of using network prediction and attention as two bridges to transfer the depth knowledge from the depth stream to the RGB stream.

898, TITLE:	Deep Fair Clustering for Visual Learning

http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Deep_Fair_Clustering_for_Visual_Learning_CVPR_2020_paper.html

AUTHORS:	Peizhao Li, Han Zhao, Hongfu Liu

HIGHLIGHT:	In light of these limitations, in this paper, we propose Deep Fair Clustering (DFC) to learn fair and clustering-favorable representations for clustering simultaneously.

899, TITLE:	Bidirectional Graph Reasoning Network for Panoptic Segmentation

http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Bidirectional_Graph_Reasoning_Network_for_Panoptic_Segmentation_CVPR_2020_paper.html

AUTHORS:	Yangxin Wu, Gengwei Zhang, Yiming Gao, Xiajun Deng, Ke Gong, Xiaodan Liang, Liang Lin

HIGHLIGHT:	We introduce a Bidirectional Graph Reasoning Network (BGRNet), which incorporates graph structure into the conventional panoptic segmentation network to mine the intra-modular and inter-modular relations within and between foreground things and background stuff classes.

900, TITLE:	Exploit Clues From Views: Self-Supervised and Regularized Learning for Multiview Object Recognition

http://openaccess.thecvf.com/content_CVPR_2020/html/Ho_Exploit_Clues_From_Views_Self-Supervised_and_Regularized_Learning_for_Multiview_CVPR_2020_paper.html

AUTHORS:	Chih-Hui Ho, Bo Liu, Tz-Ying Wu, Nuno Vasconcelos

HIGHLIGHT:	In this work, the problem of multiview self-supervised learning (MV-SSL) is investigated, where only image to object association is given.

901, TITLE:	Spherical Space Domain Adaptation With Robust Pseudo-Label Loss

http://openaccess.thecvf.com/content_CVPR_2020/html/Gu_Spherical_Space_Domain_Adaptation_With_Robust_Pseudo-Label_Loss_CVPR_2020_paper.html

AUTHORS:	Xiang Gu, Jian Sun, Zongben Xu

HIGHLIGHT:	In this paper, we propose a novel adversarial DA approach completely defined in spherical feature space, in which we define spherical classifier for label prediction and spherical domain discriminator for discriminating domain labels.

902, TITLE:	Stochastic Classifiers for Unsupervised Domain Adaptation

http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Stochastic_Classifiers_for_Unsupervised_Domain_Adaptation_CVPR_2020_paper.html

AUTHORS:	Zhihe Lu, Yongxin Yang, Xiatian Zhu, Cong Liu, Yi-Zhe Song, Tao Xiang

HIGHLIGHT:	In this paper, we introduce a novel method called STochastic clAssifieRs (STAR) for addressing this problem.

903, TITLE:	Unsupervised Learning of Intrinsic Structural Representation Points

http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Unsupervised_Learning_of_Intrinsic_Structural_Representation_Points_CVPR_2020_paper.html

AUTHORS:	Nenglun Chen, Lingjie Liu, Zhiming Cui, Runnan Chen, Duygu Ceylan, Changhe Tu, Wenping Wang

HIGHLIGHT:	We present a simple yet interpretable unsupervised method for learning a new structural representation in the form of 3D structure points.

904, TITLE:	PolyTransform: Deep Polygon Transformer for Instance Segmentation

http://openaccess.thecvf.com/content_CVPR_2020/html/Liang_PolyTransform_Deep_Polygon_Transformer_for_Instance_Segmentation_CVPR_2020_paper.html

AUTHORS:	Justin Liang, Namdar Homayounfar, Wei-Chiu Ma, Yuwen Xiong, Rui Hu, Raquel Urtasun

HIGHLIGHT:	In this paper, we propose PolyTransform, a novel instance segmentation algorithm that produces precise, geometry-preserving masks by combining the strengths of prevailing segmentation approaches and modern polygon-based methods.

905, TITLE:	Interactive Two-Stream Decoder for Accurate and Fast Saliency Detection

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Interactive_Two-Stream_Decoder_for_Accurate_and_Fast_Saliency_Detection_CVPR_2020_paper.html

AUTHORS:	Huajun Zhou, Xiaohua Xie, Jian-Huang Lai, Zixuan Chen, Lingxiao Yang

HIGHLIGHT:	In this paper, we first analyze such correlation and then propose an interactive two-stream decoder to explore multiple cues, including saliency, contour and their correlation.

906, TITLE:	Towards Better Generalization: Joint Depth-Pose Learning Without PoseNet

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Towards_Better_Generalization_Joint_Depth-Pose_Learning_Without_PoseNet_CVPR_2020_paper.html

AUTHORS:	Wang Zhao, Shaohui Liu, Yezhi Shu, Yong-Jin Liu

HIGHLIGHT:	In this work, we tackle the essential problem of scale inconsistency for self supervised joint depth-pose learning.

907, TITLE: LT-Net: Label Transfer by Learning Reversible Voxel-Wise Correspondence for One-Shot Medical Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_LT-Net_Label_Transfer_by_Learning_Reversible_Voxel-Wise_Correspondence_for_One-Shot_CVPR_2020_paper.html
AUTHORS: Shuxin Wang, Shilei Cao, Dong Wei, Renzhen Wang, Kai Ma, Liansheng Wang, Deyu Meng, Yefeng Zheng
HIGHLIGHT: We introduce a one-shot segmentation method to alleviate the burden of manual annotation for medical images.

908, TITLE: FGN: Fully Guided Network for Few-Shot Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_FGN_Fully_Guided_Network_for_Few-Shot_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Zhibo Fan, Jin-Gang Yu, Zhihao Liang, Jiarong Ou, Changxin Gao, Gui-Song Xia, Yuanqing Li
HIGHLIGHT: This paper presents a Fully Guided Network (FGN) for few-shot instance segmentation.

909, TITLE: A Quantum Computational Approach to Correspondence Problems on Point Sets
http://openaccess.thecvf.com/content_CVPR_2020/html/Golyanik_A_Quantum_Computational_Approach_to_Correspondence_Problems_on_Point_Sets_CVPR_2020_paper.html
AUTHORS: Vladislav Golyanik, Christian Theobalt
HIGHLIGHT: We review AQC and derive a new algorithm for correspondence problems on point sets suitable for execution on AQC.

910, TITLE: Data-Efficient Semi-Supervised Learning by Reliable Edge Mining
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Data-Efficient_Semi-Supervised_Learning_by_Reliable_Edge_Mining_CVPR_2020_paper.html
AUTHORS: Peibin Chen, Tao Ma, Xu Qin, Weidi Xu, Shuchang Zhou
HIGHLIGHT: We propose Reliable Edge Mining (REM), which forms a reliable graph by only selecting reliable and useful edges.

911, TITLE: NestedVAE: Isolating Common Factors via Weak Supervision
http://openaccess.thecvf.com/content_CVPR_2020/html/Vowels_NestedVAE_Isolating_Common_Factors_via_Weak_Supervision_CVPR_2020_paper.html
AUTHORS: Matthew J. Vowels, Necati Cihan Camgoz, Richard Bowden
HIGHLIGHT: To isolate the common factors we combine the theory of deep latent variable models with information bottleneck theory for scenarios whereby data may be naturally paired across domains and no additional supervision is required.

912, TITLE: Progressive Adversarial Networks for Fine-Grained Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Progressive_Adversarial_Networks_for_Fine-Grained_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Sinan Wang, Xinyang Chen, Yunbo Wang, Mingsheng Long, Jianmin Wang
HIGHLIGHT: This paper presents the Progressive Adversarial Networks (PAN) to align fine-grained categories across domains with a curriculum-based adversarial learning framework.

913, TITLE: A Disentangling Invertible Interpretation Network for Explaining Latent Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Esser_A_Disentangling_Invertible_Interpretation_Network_for_Explaining_Latent_Representations_CVPR_2020_paper.html
AUTHORS: Patrick Esser, Robin Rombach, Bjorn Ommer
HIGHLIGHT: We formulate interpretation as a translation of hidden representations onto semantic concepts that are comprehensible to the user.

914, TITLE: Modeling the Background for Incremental Learning in Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cermelli_Modeling_the_Background_for_Incremental_Learning_in_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo, Elisa Ricci, Barbara Caputo
HIGHLIGHT: In this work we revisit classical incremental learning methods, proposing a new distillation-based framework which explicitly accounts for this shift.

915, TITLE: Interpreting the Latent Space of GANs for Semantic Face Editing
http://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Interpreting_the_Latent_Space_of_GANs_for_Semantic_Face_Editing_CVPR_2020_paper.html
AUTHORS: Yujun Shen, Jinjin Gu, Xiaoou Tang, Bolei Zhou

HIGHLIGHT:	In this work, we propose a novel framework, called InterFaceGAN, for semantic face editing by interpreting the latent semantics learned by GANs.

916, TITLE:	Super-BPD: Super Boundary-to-Pixel Direction for Fast Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wan_Super-BPD_Super_Boundary-to-Pixel_Direction_for_Fast_Image_Segmentation_CVPR_2020_paper.html
AUTHORS:	Jianqiang Wan, Yang Liu, Donglai Wei, Xiang Bai, Yongchao Xu
HIGHLIGHT:	In this paper, we propose a fast image segmentation method based on a novel super boundary-to-pixel direction (super-BPD) and a customized segmentation algorithm with super-BPD.

917, TITLE:	Self-Learning With Rectification Strategy for Human Parsing
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Self-Learning_With_Rectification_Strategy_for_Human_Parsing_CVPR_2020_paper.html
AUTHORS:	Tao Li, Zhiyuan Liang, Sanyuan Zhao, Jiahao Gong, Jianbing Shen
HIGHLIGHT:	In this paper, we solve the sample shortage problem in the human parsing task.

918, TITLE:	Hyperbolic Visual Embedding Learning for Zero-Shot Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Hyperbolic_Visual_Embedding_Learning_for_Zero-Shot_Recognition_CVPR_2020_paper.html
AUTHORS:	Shaoteng Liu, Jingjing Chen, Liangming Pan, Chong-Wah Ngo, Tat-Seng Chua, Yu-Gang Jiang
HIGHLIGHT:	This paper proposes a Hyperbolic Visual Embedding Learning Network for zero-shot recognition.

919, TITLE:	Sequential Mastery of Multiple Visual Tasks: Networks Naturally Learn to Learn and Forget to Forget
http://openaccess.thecvf.com/content_CVPR_2020/html/Davidson_Sequential_Mastery_of_Multiple_Visual_Tasks_Networks_Naturally_Learn_to_CVPR_2020_paper.html
AUTHORS:	Guy Davidson, Michael C. Mozer
HIGHLIGHT:	We explore the behavior of a standard convolutional neural net in a continual-learning setting that introduces visual classification tasks sequentially and requires the net to master new tasks while preserving mastery of previously learned tasks.

920, TITLE:	Distilling Effective Supervision From Severe Label Noise
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Distilling_Effective_Supervision_From_Severe_Label_Noise_CVPR_2020_paper.html
AUTHORS:	Zizhao Zhang, Han Zhang, Sercan O. Arik, Honglak Lee, Tomas Pfister
HIGHLIGHT:	We present a holistic framework to train deep neural networks in a way that is highly invulnerable to label noise.

921, TITLE:	Eternal Sunshine of the Spotless Net: Selective Forgetting in Deep Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Golatkar_Eternal_Sunshine_of_the_Spotless_Net_Selective_Forgetting_in_Deep_CVPR_2020_paper.html
AUTHORS:	Aditya Golatkar, Alessandro Achille, Stefano Soatto
HIGHLIGHT:	We propose a method for "scrubbing" the weights clean of information about a particular set of training data.

922, TITLE:	CenterMask: Single Shot Instance Segmentation With Point Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_CenterMask_Single_Shot_Instance_Segmentation_With_Point_Representation_CVPR_2020_paper.html
AUTHORS:	Yuqing Wang, Zhaoliang Xu, Hao Shen, Baoshan Cheng, Lirong Yang
HIGHLIGHT:	In this paper, we propose a single-shot instance segmentation method, which is simple, fast and accurate.

923, TITLE:	Mitigating Bias in Face Recognition Using Skewness-Aware Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Mitigating_Bias_in_Face_Recognition_Using_Skewness-Aware_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS:	Mei Wang, Weihong Deng
HIGHLIGHT:	A reinforcement learning based race balance network (RL-RBN) is proposed.

924, TITLE:	MineGAN: Effective Knowledge Transfer From GANs to Target Domains With Few Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_MineGAN_Effective_Knowledge_Transfer_From_GANs_to_Target_Domains_With_CVPR_2020_paper.html
AUTHORS:	Yaxing Wang, Abel Gonzalez-Garcia, David Berga, Luis Herranz, Fahad Shahbaz Khan, Joost van de Weijer
HIGHLIGHT:	We propose a novel knowledge transfer method for generative models based on mining the knowledge that is most beneficial to a specific target domain, either from a single or multiple pretrained GANs.

925, TITLE: DLWL: Improving Detection for Lowshot Classes With Weakly Labelled Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Ramanathan_DLWL_Improving_Detection_for_Lowshot_Classes_With_We
akly_Labelled_Data_CVPR_2020_paper.html
AUTHORS: Vignesh Ramanathan, Rui Wang, Dhruv Mahajan
HIGHLIGHT: Towards this end, we propose a modification to the FRCNN model to automatically infer label assignment for
objects proposals from weakly labelled images during training.

926, TITLE: Unsupervised Deep Shape Descriptor With Point Distribution Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_Unsupervised_Deep_Shape_Descriptor_With_Point_Distribution_Learni
ng_CVPR_2020_paper.html
AUTHORS: Yi Shi, Mengchen Xu, Shuaihang Yuan, Yi Fang
HIGHLIGHT: This paper proposes a novel probabilistic framework for the learning of unsupervised deep shape descriptors
with point distribution learning.

927, TITLE: Stylization-Based Architecture for Fast Deep Exemplar Colorization
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Stylization-
Based_Architecture_for_Fast_Deep_Exemplar_Colorization_CVPR_2020_paper.html
AUTHORS: Zhongyou Xu, Tingting Wang, Faming Fang, Yun Sheng, Guixu Zhang
HIGHLIGHT: To tackle these problems, we propose a deep exemplar colorization architecture inspired by the characteristics
of stylization in feature extracting and blending.

928, TITLE: Cars Can't Fly Up in the Sky: Improving Urban-Scene Segmentation via Height-Driven Attention Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Choi_Cars_Cant_Fly_Up_in_the_Sky_Improving_Urban-
Scene_Segmentation_CVPR_2020_paper.html
AUTHORS: Sungha Choi, Joanne T. Kim, Jaegul Choo
HIGHLIGHT: This paper exploits the intrinsic features of urban-scene images and proposes a general add-on module, called
height-driven attention networks (HANet), for improving semantic segmentation for urban-scene images.

929, TITLE: State-Aware Tracker for Real-Time Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_State-Aware_Tracker_for_Real-
Time_Video_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Xi Chen, Zuoxin Li, Ye Yuan, Gang Yu, Jianxin Shen, Donglian Qi
HIGHLIGHT: In this work, we address the task of semi-supervised video object segmentation (VOS) and explore how to make
efficient use of video property to tackle the challenge of semi-supervision.

930, TITLE: Iteratively-Refined Interactive 3D Medical Image Segmentation With Multi-Agent Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Liao_Iteratively-
Refined_Interactive_3D_Medical_Image_Segmentation_With_Multi-Agent_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS: Xuan Liao, Wenhao Li, Qisen Xu, Xiangfeng Wang, Bo Jin, Xiaoyun Zhang, Yanfeng Wang, Ya Zhang
HIGHLIGHT: We here propose to model the dynamic process of iterative interactive image segmentation as a Markov
decision process (MDP) and solve it with reinforcement learning (RL).

931, TITLE: ENSEI: Efficient Secure Inference via Frequency-Domain Homomorphic Convolution for Privacy-Preserving
Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Bian_ENSEI_Efficient_Secure_Inference_via_Frequency-
Domain_Homomorphic_Convolution_for_Privacy-Preserving_CVPR_2020_paper.html
AUTHORS: Song Bian, Tianchen Wang, Masayuki Hiromoto, Yiyu Shi, Takashi Sato
HIGHLIGHT: In this work, we propose ENSEI, a secure inference (SI) framework based on the frequency-domain secure
convolution (FDSC) protocol for the efficient execution of image inference in the encrypted domain.

932, TITLE: Multi-Scale Interactive Network for Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Pang_Multi-
Scale_Interactive_Network_for_Salient_Object_Detection_CVPR_2020_paper.html
AUTHORS: Youwei Pang, Xiaoqi Zhao, Lihe Zhang, Huchuan Lu
HIGHLIGHT: In this paper, we propose the aggregate interaction modules to integrate the features from adjacent levels, in
which less noise is introduced because of only using small up-/down-sampling rates.

933, TITLE: Interactive Multi-Label CNN Learning With Partial Labels
http://openaccess.thecvf.com/content_CVPR_2020/html/Huynh_Interactive_Multi-
Label_CNN_Learning_With_Partial_Labels_CVPR_2020_paper.html

AUTHORS: Dat Huynh, Ehsan Elhamifar
HIGHLIGHT: We introduce a new loss function that regularizes the cross-entropy loss with a cost function that measures the smoothness of labels and features of images on the data manifold.

934, TITLE: ViewAL: Active Learning With Viewpoint Entropy for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Siddiqui_ViewAL_Active_Learning_With_Viewpoint_Entropy_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Yawar Siddiqui, Julien Valentin, Matthias Niessner
HIGHLIGHT: We propose ViewAL, a novel active learning strategy for semantic segmentation that exploits viewpoint consistency in multi-view datasets.

935, TITLE: Scene-Adaptive Video Frame Interpolation via Meta-Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Choi_Scene-Adaptive_Video_Frame_Interpolation_via_Meta-Learning_CVPR_2020_paper.html
AUTHORS: Myungsub Choi, Janghoon Choi, Sungyong Baik, Tae Hyun Kim, Kyoung Mu Lee
HIGHLIGHT: In this work, we propose to adapt the model to each video by making use of additional information that is readily available at test time and yet has not been exploited in previous works.

936, TITLE: Action Segmentation With Joint Self-Supervised Temporal Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Action_Segmentation_With_Joint_Self-Supervised_Temporal_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan AlRegib, Zsolt Kira
HIGHLIGHT: To reduce the discrepancy, we propose SelfSupervised Temporal Domain Adaptation (SSTDA), which contains two self-supervised auxiliary tasks (binary and sequential domain prediction) to jointly align cross-domain feature spaces embedded with local and global temporal dynamics, achieving better performance than other Domain Adaptation (DA) approaches.

937, TITLE: Pixel Consensus Voting for Panoptic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Pixel_Consensus_Voting_for_Panoptic_Segmentation_CVPR_2020_paper.html
AUTHORS: Haochen Wang, Ruotian Luo, Michael Maire, Greg Shakhnarovich
HIGHLIGHT: The core of our approach, Pixel Consensus Voting, is a framework for instance segmentation based on the generalized Hough transform.

938, TITLE: Minimizing Discrete Total Curvature for Image Processing
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhong_Minimizing_Discrete_Total_Curvature_for_Image_Processing_CVPR_2020_paper.html
AUTHORS: Qiuxiang Zhong, Yutong Li, Yijie Yang, Yuping Duan
HIGHLIGHT: In this paper, we propose a novel curvature regularity, the total curvature (TC), by minimizing the normal curvatures along different directions.

939, TITLE: Towards Robust Image Classification Using Sequential Attention Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Zoran_Towards_Robust_Image_Classification_Using_Sequential_Attention_Models_CVPR_2020_paper.html
AUTHORS: Daniel Zoran, Mike Chrzanowski, Po-Sen Huang, Sven Gowal, Alex Mott, Pushmeet Kohli
HIGHLIGHT: In this paper we propose to augment a modern neural-network architecture with an attention model inspired by human perception.

940, TITLE: Discovering Synchronized Subsets of Sequences: A Large Scale Solution
http://openaccess.thecvf.com/content_CVPR_2020/html/Sariyanidi_Discovering_Synchronized_Subsets_of_Sequences_A_Large_Scale_Solution_CVPR_2020_paper.html
AUTHORS: Evangelos Sariyanidi, Casey J. Zampella, Keith G. Bartley, John D. Herrington, Theodore D. Satterthwaite, Robert T. Schultz, Birkan Tunc
HIGHLIGHT: We present an approximate, but highly efficient and scalable, method that represents the search space as a union of sets called epsilon-expanded clusters, one of which is theoretically guaranteed to contain the largest subset of synchronized sequences.

941, TITLE: Going Deeper With Lean Point Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Le_Going_Deeper_With_Lean_Point_Networks_CVPR_2020_paper.html
AUTHORS: Eric-Tuan Le, Iasonas Kokkinos, Niloy J. Mitra
HIGHLIGHT: In this work we introduce Lean Point Networks (LPNs) to train deeper and more accurate point processing networks by relying on three novel point processing blocks that improve memory consumption, inference time, and accuracy: a convolution-type block for point sets that blends neighborhood information in a memory-efficient manner; a crosslink block that

efficiently shares information across low- and high-resolution processing branches; and a multi-resolution point cloud processing block for faster diffusion of information.

942, TITLE: Efficient and Robust Shape Correspondence via Sparsity-Enforced Quadratic Assignment
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_Efficient_and_Robust_Shape_Correspondence_via_Sparsity-Enforced_Quadratic_Assignment_CVPR_2020_paper.html
AUTHORS: Rui Xiang, Rongjie Lai, Hongkai Zhao
HIGHLIGHT: In this work, we introduce a novel local pairwise descriptor and then develop a simple, effective iterative method to solve the resulting quadratic assignment through sparsity control for shape correspondence between two approximate isometric surfaces.

943, TITLE: Explainable Object-Induced Action Decision for Autonomous Vehicles
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Explainable_Object-Induced_Action_Decision_for_Autonomous_Vehicles_CVPR_2020_paper.html
AUTHORS: Yiran Xu, Xiaoyin Yang, Lihang Gong, Hsuan-Chu Lin, Tz-Ying Wu, Yunsheng Li, Nuno Vasconcelos
HIGHLIGHT: A new paradigm is proposed for autonomous driving. The new paradigm lies between the end-to-end and pipelined approaches, and is inspired by how humans solve the problem.

944, TITLE: Spatially Attentive Output Layer for Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Spatially_Attentive_Output_Layer_for_Image_Classification_CVPR_2020_paper.html
AUTHORS: Ildoo Kim, Woonhyuk Baek, Sungwoong Kim
HIGHLIGHT: In this paper, we propose a novel spatial output layer on top of the existing convolutional feature maps to explicitly exploit the location-specific output information.

945, TITLE: Attack to Explain Deep Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Jalwana_Attack_to_Explain_Deep_Representation_CVPR_2020_paper.html
AUTHORS: Mohammad A. A. K. Jalwana, Naveed Akhtar, Mohammed Bennamoun, Ajmal Mian
HIGHLIGHT: This paper counter-argues and proposes the first attack on deep learning that aims at explaining the learned representation instead of fooling it.

946, TITLE: Computing Valid P-Values for Image Segmentation by Selective Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Tanizaki_Computing_Valid_P-Values_for_Image_Segmentation_by_Selective_Inference_CVPR_2020_paper.html
AUTHORS: Kosuke Tanizaki, Noriaki Hashimoto, Yu Inatsu, Hidekata Hontani, Ichiro Takeuchi
HIGHLIGHT: To overcome this difficulty, we introduce a statistical approach called selective inference, and develop a framework for computing valid p-values in which segmentation bias is properly accounted for.

947, TITLE: Unsupervised Learning From Video With Deep Neural Embeddings
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhuang_Unsupervised_Learning_From_Video_With_Deep_Neural_Embeddings_CVPR_2020_paper.html
AUTHORS: Chengxu Zhuang, Tianwei She, Alex Andonian, Max Sobol Mark, Daniel Yamins
HIGHLIGHT: Here we present the Video Instance Embedding (VIE) framework, which trains deep nonlinear embeddings on video sequence inputs.

948, TITLE: Partial Weight Adaptation for Robust DNN Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_Partial_Weight_Adaptation_for_Robust_DNN_Inference_CVPR_2020_paper.html
AUTHORS: Xiufeng Xie, Kyu-Han Kim
HIGHLIGHT: We present GearNN, an adaptive inference architecture that accommodates DNN inputs with varying distortions.

949, TITLE: Probability Weighted Compact Feature for Domain Adaptive Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Probability_Weighted_Compact_Feature_for_Domain_Adaptive_Retrieval_CVPR_2020_paper.html
AUTHORS: Fuxiang Huang, Lei Zhang, Yang Yang, Xichuan Zhou
HIGHLIGHT: In this paper, considering the practical application, we focus on challenging cross-domain retrieval.

950, TITLE: Where Does It End? - Reasoning About Hidden Surfaces by Object Intersection Constraints
http://openaccess.thecvf.com/content_CVPR_2020/html/Strecke_Where_Does_It_End_-_Reasoning_About_Hidden_Surfaces_by_CVPR_2020_paper.html

AUTHORS: Michael Strecke, Jorg Stuckler
HIGHLIGHT: In this paper we propose Co-Section, an optimization-based approach to 3D dynamic scene reconstruction, which infers hidden shape information from intersection constraints.

951, TITLE: PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_PolarNet_An_Improved_Grid_Representation_for_Online_LiDAR_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, Hassan Foroosh
HIGHLIGHT: The combination of the aforementioned challenges motivates us to propose a new LiDAR-specific, KNN-free segmentation algorithm - PolarNet.

952, TITLE: Pathological Retinal Region Segmentation From OCT Images Using Geometric Relation Based Augmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Mahapatra_Pathological_Retinal_Region_Segmentation_From_OCT_Images_Using_Geometric_Relation_CVPR_2020_paper.html
AUTHORS: Dwarikanath Mahapatra, Behzad Bozorgtabar, Ling Shao
HIGHLIGHT: We propose improvements over previous GAN-based medical image synthesis methods by jointly encoding the intrinsic relationship of geometry and shape.

953, TITLE: Transferring and Regularizing Prediction for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Transferring_and_Regularizing_Prediction_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Yiheng Zhang, Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Dong Liu, Tao Mei
HIGHLIGHT: In this paper, we novelly exploit the intrinsic properties of semantic segmentation to alleviate such problem for model transfer.

954, TITLE: PREDICT & CLUSTER: Unsupervised Skeleton Based Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Su_PREDICT__CLUSTER_Unsupervised_Skeleton_Based_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Kun Su, Xiulong Liu, Eli Shlizerman
HIGHLIGHT: We propose a novel system for unsupervised skeleton-based action recognition.

955, TITLE: Model Adaptation: Unsupervised Domain Adaptation Without Source Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Model_Adaptation_Unsupervised_Domain_Adaptation_Without_Source_Data_CVPR_2020_paper.html
AUTHORS: Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, Si Wu
HIGHLIGHT: In this paper, we investigate a challenging unsupervised domain adaptation setting --- unsupervised model adaptation.

956, TITLE: Evade Deep Image Retrieval by Stashing Private Images in the Hash Space
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiao_Evade_Deep_Image_Retrieval_by_Stashing_Private_Images_in_the_CVPR_2020_paper.html
AUTHORS: Yanru Xiao, Cong Wang, Xing Gao
HIGHLIGHT: In this paper, we propose a new mechanism based on adversarial examples to "stash" private images in the deep hash space while maintaining perceptual similarity.

957, TITLE: Advisable Learning for Self-Driving Vehicles by Internalizing Observation-to-Action Rules
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Advisable_Learning_for_Self-Driving_Vehicles_by_Internalizing_Observation-to-Action_Rules_CVPR_2020_paper.html
AUTHORS: Jinkyu Kim, Suhong Moon, Anna Rohrbach, Trevor Darrell, John Canny
HIGHLIGHT: We propose a new approach that learns vehicle control with the help of human advice.

958, TITLE: ProAlignNet: Unsupervised Learning for Progressively Aligning Noisy Contours
http://openaccess.thecvf.com/content_CVPR_2020/html/Veeravasarapu_ProAlignNet_Unsupervised_Learning_for_Progressively_Aligning_Noisy_Contours_CVPR_2020_paper.html
AUTHORS: VSR Veeravasarapu, Abhishek Goel, Deepak Mittal, Maneesh Singh
HIGHLIGHT: This work presents a novel ConvNet, "ProAlignNet," that accounts for large scale misalignments and complex transformations between the contour shapes.

959, TITLE: Attribution in Scale and Space
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Attribution_in_Scale_and_Space_CVPR_2020_paper.html
AUTHORS: Shawn Xu, Subhashini Venugopalan, Mukund Sundararajan

HIGHLIGHT: We propose a new technique called Blur Integrated Gradients (Blur IG).

960, TITLE: Towards Causal VQA: Revealing and Reducing Spurious Correlations by Invariant and Covariant Semantic Editing
http://openaccess.thecvf.com/content_CVPR_2020/html/Agarwal_Towards_Causal_VQA_Revealing_and_Reducing_Spurious_Correlations_by_Invariant_CVPR_2020_paper.html
AUTHORS: Vedika Agarwal, Rakshith Shetty, Mario Fritz
HIGHLIGHT: In this paper, we propose a novel way to analyze and measure the robustness of the state of the art models w.r.t semantic visual variations as well as propose ways to make models more robust against spurious correlations.

961, TITLE: Deep Relational Reasoning Graph Network for Arbitrary Shape Text Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Deep_Relational_Reasoning_Graph_Network_for_Arbitrary_Shape_Text_Detection_CVPR_2020_paper.html
AUTHORS: Shi-Xue Zhang, Xiaobin Zhu, Jie-Bo Hou, Chang Liu, Chun Yang, Hongfa Wang, Xu-Cheng Yin
HIGHLIGHT: In this paper, we propose a novel unified relational reasoning graph network for arbitrary shape text detection.

962, TITLE: Large-Scale Object Detection in the Wild From Imbalanced Multi-Labels
http://openaccess.thecvf.com/content_CVPR_2020/html/Peng_Large-Scale_Object_Detection_in_the_Wild_From_Imbalanced_Multi-Labels_CVPR_2020_paper.html
AUTHORS: Junran Peng, Xingyuan Bu, Ming Sun, Zhaoxiang Zhang, Tieniu Tan, Junjie Yan
HIGHLIGHT: In this work, we quantitatively analyze these label problems and provide a simple but effective solution.

963, TITLE: BBN: Bilateral-Branch Network With Cumulative Learning for Long-Tailed Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_BBN_Bilateral-Branch_Network_With_Cumulative_Learning_for_Long-Tailed_Visual_Recognition_CVPR_2020_paper.html
AUTHORS: Boyan Zhou, Quan Cui, Xiu-Shen Wei, Zhao-Min Chen
HIGHLIGHT: Therefore, we propose a unified Bilateral-Branch Network (BBN) to take care of both representation learning and classifier learning simultaneously, where each branch does perform its own duty separately.

964, TITLE: Momentum Contrast for Unsupervised Visual Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/He_Momentum_Contrast_for_Unsupervised_Visual_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, Ross Girshick
HIGHLIGHT: We present Momentum Contrast (MoCo) for unsupervised visual representation learning.

965, TITLE: Classifying, Segmenting, and Tracking Object Instances in Video with Mask Propagation
http://openaccess.thecvf.com/content_CVPR_2020/html/Bertasius_Classifying_Segmenting_and_Tracking_Object_Instances_in_Video_with_Mask_CVPR_2020_paper.html
AUTHORS: Gedas Bertasius, Lorenzo Torresani
HIGHLIGHT: We introduce a method for simultaneously classifying, segmenting and tracking object instances in a video sequence.

966, TITLE: Weakly Supervised Fine-Grained Image Classification via Guassian Mixture Model Oriented Discriminative Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Weakly_Supervised_Fine-Grained_Image_Classification_via_Guassian_Mixture_Model_Oriented_CVPR_2020_paper.html
AUTHORS: Zhihui Wang, Shijie Wang, Shuhui Yang, Haojie Li, Jianjun Li, Zezhou Li
HIGHLIGHT: In this paper, we propose an end-to-end Discriminative Feature-oriented Gaussian Mixture Model (DF-GMM), to address the problem of discriminative region diffusion and find better fine-grained details.

967, TITLE: Bridging the Gap Between Anchor-Based and Anchor-Free Detection via Adaptive Training Sample Selection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Bridging_the_Gap_Between_Anchor-Based_and_Anchor-Free_Detection_via_Adaptive_CVPR_2020_paper.html
AUTHORS: Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, Stan Z. Li
HIGHLIGHT: In this paper, we first point out that the essential difference between anchor-based and anchor-free detection is actually how to define positive and negative training samples, which leads to the performance gap between them.

968, TITLE: Learning User Representations for Open Vocabulary Image Hashtag Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Durand_Learning_User_Representations_for_Open_Vocabulary_Image_Hashtag_Prediction_CVPR_2020_paper.html
AUTHORS: Thibaut Durand

HIGHLIGHT: In this paper, we introduce an open vocabulary model for image hashtag prediction - the task of mapping an image to its accompanying hashtags.

969, TITLE: Sketch Less for More: On-the-Fly Fine-Grained Sketch-Based Image Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Bhunia_Sketch_Less_for_More_On-the-Fly_Fine-Grained_Sketch-Based_Image_Retrieval_CVPR_2020_paper.html
AUTHORS: Ayan Kumar Bhunia, Yongxin Yang, Timothy M. Hospedales, Tao Xiang, Yi-Zhe Song
HIGHLIGHT: In this paper, we reformulate the conventional FG-SBIR framework to tackle these challenges, with the ultimate goal of retrieving the target photo with the least number of strokes possible.

970, TITLE: Few-Shot Pill Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Ling_Few-Shot_Pill_Recognition_CVPR_2020_paper.html
AUTHORS: Suiyi Ling, Andreas Pastor, Jing Li, Zhaohui Che, Junle Wang, Jieun Kim, Patrick Le Callet
HIGHLIGHT: In this study, a new pill image database, namely CURE, is first developed with more varied imaging conditions and instances for each pill category. Secondly, a W2-net is proposed for better pill segmentation. Thirdly, a Multi-Stream (MS) deep network that captures task-related features along with a novel two-stage training methodology are proposed.

971, TITLE: PointRend: Image Segmentation As Rendering
http://openaccess.thecvf.com/content_CVPR_2020/html/Kirillov_PointRend_Image_Segmentation_As_Rendering_CVPR_2020_paper.html
AUTHORS: Alexander Kirillov, Yuxin Wu, Kaiming He, Ross Girshick
HIGHLIGHT: We present a new method for efficient high-quality image segmentation of objects and scenes.

972, TITLE: ABCNet: Real-Time Scene Text Spotting With Adaptive Bezier-Curve Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_ABCNet_Real-Time_Scene_Text_Spotting_With_Adaptive_Bezier-Curve_Network_CVPR_2020_paper.html
AUTHORS: Yuliang Liu, Hao Chen, Chunhua Shen, Tong He, Lianwen Jin, Liangwei Wang
HIGHLIGHT: Our contributions are three-fold: 1) For the first time, we adaptively fit oriented or curved text by a parameterized Bezier curve. 2) We design a novel BezierAlign layer for extracting accurate convolution features of a text instance with arbitrary shapes, significantly improving the precision compared with previous methods.

973, TITLE: Learning Temporal Co-Attention Models for Unsupervised Video Action Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Gong_Learning_Temporal_Co-Attention_Models_for_Unsupervised_Video_Action_Localization_CVPR_2020_paper.html
AUTHORS: Guoqiang Gong, Xinghan Wang, Yadong Mu, Qi Tian
HIGHLIGHT: To solve ACL, we propose a two-step "clustering + localization" iterative procedure.

974, TITLE: Spatiotemporal Fusion in 3D CNNs: A Probabilistic View
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Spatiotemporal_Fusion_in_3D_CNNs_A_Probabilistic_View_CVPR_2020_paper.html
AUTHORS: Yizhou Zhou, Xiaoyan Sun, Chong Luo, Zheng-Jun Zha, Wenjun Zeng
HIGHLIGHT: In this paper, we propose to convert the spatiotemporal fusion strategies into a probability space, which allows us to perform network-level evaluations of various fusion strategies without having to train them separately.

975, TITLE: Uncertainty-Aware Score Distribution Learning for Action Quality Assessment
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Uncertainty-Aware_Score_Distribution_Learning_for_Action_Quality_Assessment_CVPR_2020_paper.html
AUTHORS: Yansong Tang, Zanlin Ni, Jiahuan Zhou, Danyang Zhang, Jiwen Lu, Ying Wu, Jie Zhou
HIGHLIGHT: To address this issue, we propose an uncertainty-aware score distribution learning (USDL) approach for action quality assessment (AQA).

976, TITLE: Learning Interactions and Relationships Between Movie Characters
http://openaccess.thecvf.com/content_CVPR_2020/html/Kukleva_Learning_Interactions_and_Relationships_Between_Movie_Characters_CVPR_2020_paper.html
AUTHORS: Anna Kukleva, Makarand Tapaswi, Ivan Laptev
HIGHLIGHT: In this work, we propose neural models to learn and jointly predict interactions, relationships, and the pair of characters that are involved.

977, TITLE: Video Panoptic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Video_Panoptic_Segmentation_CVPR_2020_paper.html
AUTHORS: Dahun Kim, Sanghyun Woo, Joon-Young Lee, In So Kweon

HIGHLIGHT:	In this paper, we propose and explore a new video extension of this task, called video panoptic segmentation.

978, TITLE:	Understanding Human Hands in Contact at Internet Scale
http://openaccess.thecvf.com/content_CVPR_2020/html/Shan_Understanding_Human_Hands_in_Contact_at_Internet_Scale_CVPR_2020_paper.html
AUTHORS:	Dandan Shan, Jiaqi Geng, Michelle Shu, David F. Fouhey
HIGHLIGHT:	This paper proposes steps towards this by inferring a rich representation of hands engaged in interaction method that includes: hand location, side, contact state, and a box around the object in contact.

979, TITLE:	End-to-End Learning of Visual Representations From Uncurated Instructional Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Miech_End-to-End_Learning_of_Visual_Representations_From_Uncurated_Instructional_Videos_CVPR_2020_paper.html
AUTHORS:	Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, Andrew Zisserman
HIGHLIGHT:	In this work we propose a new learning approach, MIL-NCE, capable of addressing mis- alignments inherent in narrated videos.

980, TITLE:	You2Me: Inferring Body Pose in Egocentric Video via First and Second Person Interactions
http://openaccess.thecvf.com/content_CVPR_2020/html/Ng_You2Me_Inferring_Body_Pose_in_Egocentric_Video_via_First_and_CVPR_2020_paper.html
AUTHORS:	Evonne Ng, Donglai Xiang, Hanbyul Joo, Kristen Grauman
HIGHLIGHT:	We propose a learning-based approach to estimate the camera wearer's 3D body pose from egocentric video sequences.

981, TITLE:	Learning a Weakly-Supervised Video Actor-Action Segmentation Model With a Wise Selection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Learning_a_Weakly-Supervised_Video_Actor-Action_Segmentation_Model_With_a_Wise_CVPR_2020_paper.html
AUTHORS:	Jie Chen, Zhiheng Li, Jiebo Luo, Chenliang Xu
HIGHLIGHT:	To overcome these challenges, we propose a general Weakly-Supervised framework with a Wise Selection of training samples and model evaluation criterion (WS^2).

982, TITLE:	Learning to Measure the Static Friction Coefficient in Cloth Contact
http://openaccess.thecvf.com/content_CVPR_2020/html/Rasheed_Learning_to_Measure_the_Static_Friction_Coefficient_in_Cloth_Contact_CVPR_2020_paper.html
AUTHORS:	Abdullah Haroon Rasheed, Victor Romero, Florence Bertails-Descoubes, Stefanie Wuhrer, Jean-Sebastien Franco, Arnaud Lazarus
HIGHLIGHT:	We propose a first vision-based measurement network for friction between cloth and a substrate, using a simple and repeatable video acquisition protocol.

983, TITLE:	SpeedNet: Learning the Speediness in Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Benaim_SpeedNet_Learning_the_Speediness_in_Videos_CVPR_2020_paper.html
AUTHORS:	Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T. Freeman, Michael Rubinstein, Michal Irani, Tali Dekel
HIGHLIGHT:	The core component in our approach is SpeedNet--a novel deep network trained to detect if a video is playing at normal rate, or if it is sped up.

984, TITLE:	Telling Left From Right: Learning Spatial Correspondence of Sight and Sound
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Telling_Left_From_Right_Learning_Spatial_Correspondence_of_Sight_and_CVPR_2020_paper.html
AUTHORS:	Karren Yang, Bryan Russell, Justin Salamon
HIGHLIGHT:	We propose a novel self-supervised task to leverage an orthogonal principle: matching spatial information in the audio stream to the positions of sound sources in the visual stream.

985, TITLE:	Visual-Textual Capsule Routing for Text-Based Video Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/McIntosh_Visual-Textual_Capsule_Routing_for_Text-Based_Video_Segmentation_CVPR_2020_paper.html
AUTHORS:	Bruce McIntosh, Kevin Duarte, Yogesh S Rawat, Mubarak Shah
HIGHLIGHT:	In this work, we focus on integration of video and text for the task of actor and action video segmentation from a sentence.

986, TITLE:	Graph-Structured Referring Expression Reasoning in the Wild

http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Graph-Structured_Referring_Expression_Reasoning_in_the_Wild_CVPR_2020_paper.html
AUTHORS: Sibei Yang, Guanbin Li, Yizhou Yu
HIGHLIGHT: In this paper, we propose a scene graph guided modular network (SGMN), which performs reasoning over a semantic graph and a scene graph with neural modules under the guidance of the linguistic structure of the expression.

987, TITLE: Say As You Wish: Fine-Grained Control of Image Caption Generation With Abstract Scene Graphs
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Say_As_You_Wish_Fine-Grained_Control_of_Image_Caption_Generation_CVPR_2020_paper.html
AUTHORS: Shizhe Chen, Qin Jin, Peng Wang, Qi Wu
HIGHLIGHT: In this work, we propose the Abstract Scene Graph (ASG) structure to represent user intention in fine-grained level and control what and how detailed the generated description should be.

988, TITLE: Hierarchical Conditional Relation Networks for Video Question Answering
http://openaccess.thecvf.com/content_CVPR_2020/html/Le_Hierarchical_Conditional_Relation_Networks_for_Video_Question_Answering_CVPR_2020_paper.html
AUTHORS: Thao Minh Le, Vuong Le, Svetha Venkatesh, Truyen Tran
HIGHLIGHT: We introduce a general-purpose reusable neural unit called Conditional Relation Network (CRN) that serves as a building block to construct more sophisticated structures for representation and reasoning over video.

989, TITLE: REVERIE: Remote Embodied Visual Referring Expression in Real Indoor Environments
http://openaccess.thecvf.com/content_CVPR_2020/html/Qi_REVERIE_Remote_Embodied_Visual_Referring_Expression_in_Real_Indoor_Environments_CVPR_2020_paper.html
AUTHORS: Yuankai Qi, Qi Wu, Peter Anderson, Xin Wang, William Yang Wang, Chunhua Shen, Anton van den Hengel
HIGHLIGHT: In the hope that it might drive progress towards more flexible and powerful human interactions with robots, we propose a dataset of varied and complex robot tasks, described in natural language, in terms of objects visible in a large set of real images.

990, TITLE: Iterative Answer Prediction With Pointer-Augmented Multimodal Transformers for TextVQA
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Iterative_Answer_Prediction_With_Pointer-Augmented_Multimodal_Transformers_for_TextVQA_CVPR_2020_paper.html
AUTHORS: Ronghang Hu, Amanpreet Singh, Trevor Darrell, Marcus Rohrbach
HIGHLIGHT: In this work, we propose a novel model for the TextVQA task based on a multimodal transformer architecture accompanied by a rich representation for text in images.

991, TITLE: SQuINTing at VQA Models: Introspecting VQA Models With Sub-Questions
http://openaccess.thecvf.com/content_CVPR_2020/html/Selvaraju_SQuINTing_at_VQA_Models_Introspecting_VQA_Models_With_Sub-Questions_CVPR_2020_paper.html
AUTHORS: Ramprasaath R. Selvaraju, Purva Tendulkar, Devi Parikh, Eric Horvitz, Marco Tulio Ribeiro, Besmira Nushi, Ece Kamar
HIGHLIGHT: To address this shortcoming, we propose an approach called Sub-Question Importance-aware Network Tuning (SQuINT), which encourages the model to attend to the same parts of the image when answering the reasoning question and the perception sub question.

992, TITLE: Vision-Language Navigation With Self-Supervised Auxiliary Reasoning Tasks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Vision-Language_Navigation_With_Self-Supervised_Auxiliary_Reasoning_Tasks_CVPR_2020_paper.html
AUTHORS: Fengda Zhu, Yi Zhu, Xiaojun Chang, Xiaodan Liang
HIGHLIGHT: In this paper, we introduce Auxiliary Reasoning Navigation (AuxRN), a framework with four self-supervised auxiliary reasoning tasks to exploit the additional training signals derived from these semantic information.

993, TITLE: Sign Language Transformers: Joint End-to-End Sign Language Recognition and Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Camgoz_Sign_Language_Transformers_Joint_End-to-End_Sign_Language_Recognition_and_Translation_CVPR_2020_paper.html
AUTHORS: Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, Richard Bowden
HIGHLIGHT: We introduce a novel transformer based architecture that jointly learns Continuous Sign Language Recognition and Translation while being trainable in an end-to-end manner.

994, TITLE: Multi-Task Collaborative Network for Joint Referring Expression Comprehension and Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Multi-Task_Collaborative_Network_for_Joint_Referring_Expression_Comprehension_and_Segmentation_CVPR_2020_paper.html
AUTHORS: Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, Rongrong Ji

HIGHLIGHT: In this paper, we propose a novel Multi-task Collaborative Network (MCN) to achieve a joint learning of REC and RES for the first time.

995, TITLE: Counterfactual Vision and Language Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Abbasnejad_Counterfactual_Vision_and_Language_Learning_CVPR_2020_paper.html
AUTHORS: Ehsan Abbasnejad, Damien Teney, Amin Parvaneh, Javen Shi, Anton van den Hengel
HIGHLIGHT: We propose a method that addresses this problem by introducing counterfactuals in the training.

996, TITLE: Iterative Context-Aware Graph Inference for Visual Dialog
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Iterative_Context-Aware_Graph_Inference_for_Visual_Dialog_CVPR_2020_paper.html
AUTHORS: Dan Guo, Hui Wang, Hanwang Zhang, Zheng-Jun Zha, Meng Wang
HIGHLIGHT: To this end, we propose a novel Context-Aware Graph (CAG) neural network.

997, TITLE: TA-Student VQA: Multi-Agents Training by Self-Questioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Xiong_TA-Student_VQA_Multi-Agents_Training_by_Self-Questioning_CVPR_2020_paper.html
AUTHORS: Peixi Xiong, Ying Wu
HIGHLIGHT: We introduce our self-questioning model with multi-agent training: TA-student VQA.

998, TITLE: Exploring Self-Attention for Image Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Exploring_Self-Attention_for_Image_Recognition_CVPR_2020_paper.html
AUTHORS: Hengshuang Zhao, Jiaya Jia, Vladlen Koltun
HIGHLIGHT: We explore variations of self-attention and assess their effectiveness for image recognition.

999, TITLE: Cops-Ref: A New Dataset and Task on Compositional Referring Expression Comprehension
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Cops-Ref_A_New_Dataset_and_Task_on_Compositional_Referring_Expression_CVPR_2020_paper.html
AUTHORS: Zhenfang Chen, Peng Wang, Lin Ma, Kwan-Yee K. Wong, Qi Wu
HIGHLIGHT: To bridge the gap, we propose a new dataset for visual reasoning in context of referring expression comprehension with two main features.

1000, TITLE: Improving Convolutional Networks With Self-Calibrated Convolutions
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Improving_Convolutional_Networks_With_Self-Calibrated_Convolutions_CVPR_2020_paper.html
AUTHORS: Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Changhu Wang, Jiashi Feng
HIGHLIGHT: In this paper, we consider how to improve the basic convolutional feature transformation process of CNNs without tuning the model architectures.

1001, TITLE: Modality Shifting Attention Network for Multi-Modal Video Question Answering
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Modality_Shifting_Attention_Network_for_Multi-Modal_Video_Question_Answering_CVPR_2020_paper.html
AUTHORS: Junyeong Kim, Minuk Ma, Trung Pham, Kyungsu Kim, Chang D. Yoo
HIGHLIGHT: This paper considers a network referred to as Modality Shifting Attention Network (MSAN) for Multimodal Video Question Answering (MVQA) task.

1002, TITLE: Learning to Structure an Image With Few Colors
http://openaccess.thecvf.com/content_CVPR_2020/html/Hou_Learning_to_Structure_an_Image_With_Few_Colors_CVPR_2020_paper.html
AUTHORS: Yunzhong Hou, Liang Zheng, Stephen Gould
HIGHLIGHT: To this end, we propose a color quantization network, ColorCNN, which learns to structure the images from the classification loss in an end-to-end manner.

1003, TITLE: On the General Value of Evidence, and Bilingual Scene-Text Visual Question Answering
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_On_the_General_Value_of_Evidence_and_Bilingual_Scene-Text_Visual_CVPR_2020_paper.html
AUTHORS: Xinyu Wang, Yuliang Liu, Chunhua Shen, Chun Chet Ng, Canjie Luo, Lianwen Jin, Chee Seng Chan, Anton van den Hengel, Liangwei Wang

HIGHLIGHT: We present a dataset that takes a step towards addressing this problem in that it contains questions expressed in two languages, and an evaluation process that co-opts a well understood image-based metric to reflect the method's ability to reason.

1004, TITLE: From Paris to Berlin: Discovering Fashion Style Influences Around the World
http://openaccess.thecvf.com/content_CVPR_2020/html/Al-Halah_From_Paris_to_Berlin_Discovering_Fashion_Style_Influences_Around_the_CVPR_2020_paper.html
AUTHORS: Ziad Al-Halah, Kristen Grauman
HIGHLIGHT: We introduce an approach that detects which cities influence which other cities in terms of propagating their styles.

1005, TITLE: A Local-to-Global Approach to Multi-Modal Movie Scene Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Rao_A_Local-to-Global_Approach_to_Multi-Modal_Movie_Scene_Segmentation_CVPR_2020_paper.html
AUTHORS: Anyi Rao, Linning Xu, Yu Xiong, Guodong Xu, Qingqiu Huang, Bolei Zhou, Dahua Lin
HIGHLIGHT: Towards this goal, we scale up the scene segmentation task by building a large-scale video dataset MovieScenes, which contains 21K annotated scene segments from 150 movies. We further propose a local-to-global scene segmentation framework, which integrates multi-modal information across three levels, i.e. clip, segment, and movie.

1006, TITLE: G-TAD: Sub-Graph Localization for Temporal Action Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_G-TAD_Sub-Graph_Localization_for_Temporal_Action_Detection_CVPR_2020_paper.html
AUTHORS: Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, Bernard Ghanem
HIGHLIGHT: In this work, we propose a graph convolutional network (GCN) model to adaptively incorporate multi-level semantic context into video features and cast temporal action detection as a sub-graph localization problem.

1007, TITLE: Detailed 2D-3D Joint Representation for Human-Object Interaction
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Detailed_2D-3D_Joint_Representation_for_Human-Object_Interaction_CVPR_2020_paper.html
AUTHORS: Yong-Lu Li, Xinpeng Liu, Han Lu, Shiyi Wang, Junqi Liu, Jiefeng Li, Cewu Lu
HIGHLIGHT: In light of these, we propose a detailed 2D-3D joint representation learning method.
To better evaluate the 2D ambiguity processing capacity of models, we propose a new benchmark named Ambiguous-HOI consisting of hard ambiguous images.

1008, TITLE: One-Shot Adversarial Attacks on Visual Tracking With Dual Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_One-Shot_Adversarial_Attacks_on_Visual_Tracking_With_Dual_Attention_CVPR_2020_paper.html
AUTHORS: Xuesong Chen, Xiyu Yan, Feng Zheng, Yong Jiang, Shu-Tao Xia, Yong Zhao, Rongrong Ji
HIGHLIGHT: In this paper, we propose a novel one-shot adversarial attack method to generate adversarial examples for free-model single object tracking, where merely adding slight perturbations on the target patch in the initial frame causes state-of-the-art trackers to lose the target in subsequent frames.

1009, TITLE: Rethinking Classification and Localization for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Rethinking_Classification_and_Localization_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Yue Wu, Yinpeng Chen, Lu Yuan, Zicheng Liu, Lijuan Wang, Hongzhi Li, Yun Fu
HIGHLIGHT: Based upon these findings, we propose a Double-Head method, which has a fully connected head focusing on classification and a convolution head for bounding box regression.

1010, TITLE: Correspondence Networks With Adaptive Neighbourhood Consensus
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Correspondence_Networks_With_Adaptive_Neighbourhood_Consensus_CVPR_2020_paper.html
AUTHORS: Shuda Li, Kai Han, Theo W. Costain, Henry Howard-Jenkins, Victor Prisacariu
HIGHLIGHT: In this paper, we tackle the task of establishing dense visual correspondences between images containing objects of the same category.

1011, TITLE: Multiple Anchor Learning for Visual Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Ke_Multiple_Anchor_Learning_for_Visual_Object_Detection_CVPR_2020_paper.html
AUTHORS: Wei Ke, Tianliang Zhang, Zeyi Huang, Qixiang Ye, Jianzhuang Liu, Dong Huang
HIGHLIGHT: In this paper, we propose a Multiple Instance Learning (MIL) approach that selects anchors and jointly optimizes the two modules of a CNN-based object detector.

1012, TITLE:	PhraseCut: Language-Based Image Segmentation in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_PhraseCut_Language-Based_Image_Segmentation_in_the_Wild_CVPR_2020_paper.html
AUTHORS:	Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, Subhransu Maji
HIGHLIGHT:	We consider the problem of segmenting image regions given a natural language phrase, and study it on a novel dataset of 77,262 images and 345,486 phrase-region pairs.

1013, TITLE:	Mask Encoding for Single Shot Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Mask_Encoding_for_Single_Shot_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS:	Rufeng Zhang, Zhi Tian, Chunhua Shen, Mingyu You, Youliang Yan
HIGHLIGHT:	In this work, we propose a simple single-shot instance segmentation framework, termed mask encoding based instance segmentation (MEInst).

1014, TITLE:	Action Genome: Actions As Compositions of Spatio-Temporal Scene Graphs
http://openaccess.thecvf.com/content_CVPR_2020/html/Ji_Action_Genome_Actions_As_Compositions_of_Spatio-Temporal_Scene_Graphs_CVPR_2020_paper.html
AUTHORS:	Jingwei Ji, Ranjay Krishna, Li Fei-Fei, Juan Carlos Niebles
HIGHLIGHT:	Inspired by evidence that the prototypical unit of an event is an action-object interaction, we introduce Action Genome, a representation that decomposes actions into spatio-temporal scene graphs.

1015, TITLE:	Learning Unseen Concepts via Hierarchical Decomposition and Composition
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_Unseen_Concepts_via_Hierarchical_Decomposition_and_Composition_CVPR_2020_paper.html
AUTHORS:	Muli Yang, Cheng Deng, Junchi Yan, Xianglong Liu, Dacheng Tao
HIGHLIGHT:	We propose to learn unseen concepts in a hierarchical decomposition-and-composition manner.

1016, TITLE:	Hi-CMD: Hierarchical Cross-Modality Disentanglement for Visible-Infrared Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Choi_Hi-CMD_Hierarchical_Cross-Modality_Disentanglement_for_Visible-Infrared_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS:	Seokeon Choi, Sumin Lee, Youngeun Kim, Taekyung Kim, Changick Kim
HIGHLIGHT:	To reduce both intra- and cross-modality discrepancies, we propose a Hierarchical Cross-Modality Disentanglement (Hi-CMD) method, which automatically disentangles ID-discriminative factors and ID-excluded factors from visible-thermal images.

1017, TITLE:	In Defense of Grid Features for Visual Question Answering
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_In_Defense_of_Grid_Features_for_Visual_Question_Answering_CVPR_2020_paper.html
AUTHORS:	Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, Xinlei Chen
HIGHLIGHT:	In this paper, we revisit grid features for VQA, and find they can work surprisingly well -- running more than an order of magnitude faster with the same accuracy (e.g. if pre-trained in a similar fashion).

1018, TITLE:	Multi-Mutual Consistency Induced Transfer Subspace Learning for Human Motion Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Multi-Mutual_Consistency_Induced_Transfer_Subspace_Learning_for_Human_Motion_Segmentation_CVPR_2020_paper.html
AUTHORS:	Tao Zhou, Huazhu Fu, Chen Gong, Jianbing Shen, Ling Shao, Fatih Porikli
HIGHLIGHT:	To this end, we propose a novel multi-mutual consistency induced transfer subspace learning framework for human motion segmentation.

1019, TITLE:	Dense Regression Network for Video Grounding
http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_Dense_Regression_Network_for_Video_Grounding_CVPR_2020_paper.html
AUTHORS:	Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Mingkui Tan, Chuang Gan
HIGHLIGHT:	The key idea of this paper is to use the distances between the frame within the ground truth and the starting (ending) frame as dense supervisions to improve the video grounding accuracy.

1020, TITLE:	Neural Architecture Search for Lightweight Non-Local Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Neural_Architecture_Search_for_Lightweight_Non-Local_Networks_CVPR_2020_paper.html
AUTHORS:	Yingwei Li, Xiaojie Jin, Jieru Mei, Xiaochen Lian, Linjie Yang, Cihang Xie, Qihang Yu, Yuyin Zhou, Song Bai, Alan L. Yuille

HIGHLIGHT: We propose AutoNL to overcome the above two obstacles.

1021, TITLE: Learning Saliency Propagation for Semi-Supervised Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Learning_Saliency_Propagation_for_Semi-Supervised_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, Fisher Yu
HIGHLIGHT: We propose ShapeProp, which learns to activate the salient regions within the object detection and propagate the areas to the whole instance through an iterative learnable message passing module.

1022, TITLE: Speech2Action: Cross-Modal Supervision for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Nagrani_Speech2Action_Cross-Modal_Supervision_for_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Arsha Nagrani, Chen Sun, David Ross, Rahul Sukthankar, Cordelia Schmid, Andrew Zisserman
HIGHLIGHT: In this work we investigate the link between spoken words and actions in movies.

1023, TITLE: Normalized and Geometry-Aware Self-Attention Network for Image Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Normalized_and_Geometry-Aware_Self-Attention_Network_for_Image_Captioning_CVPR_2020_paper.html
AUTHORS: Longteng Guo, Jing Liu, Xinxin Zhu, Peng Yao, Shichen Lu, Hanqing Lu
HIGHLIGHT: In this paper, we improve SA from two aspects to promote the performance of image captioning.

1024, TITLE: Memory Enhanced Global-Local Aggregation for Video Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Memory_Enhanced_Global-Local_Aggregation_for_Video_Object_Detection_CVPR_2020_paper.html
AUTHORS: Yihong Chen, Yue Cao, Han Hu, Liwei Wang
HIGHLIGHT: In this paper we introduce memory enhanced global-local aggregation (MEGA) network, which is among the first trials that takes full consideration of both global and local information.

1025, TITLE: Solving Mixed-Modal Jigsaw Puzzle for Fine-Grained Sketch-Based Image Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Pang_Solving_Mixed-Modal_Jigsaw_Puzzle_for_Fine-Grained_Sketch-Based_Image_Retrieval_CVPR_2020_paper.html
AUTHORS: Kaiyue Pang, Yongxin Yang, Timothy M. Hospedales, Tao Xiang, Yi-Zhe Song
HIGHLIGHT: In this paper, we propose a self-supervised alternative for representation pre-training.

1026, TITLE: LG-GAN: Label Guided Adversarial Network for Flexible Targeted Attack of Point Cloud Based Deep Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_LG-GAN_Label_Guided_Adversarial_Network_for_Flexible_Targeted_Attack_of_CVPR_2020_paper.html
AUTHORS: Hang Zhou, Dongdong Chen, Jing Liao, Kejiang Chen, Xiaoyi Dong, Kunlin Liu, Weiming Zhang, Gang Hua, Nenghai Yu
HIGHLIGHT: To overcome these shortcomings, this paper proposes a novel label guided adversarial network (LG-GAN) for real-time flexible targeted point cloud attack.

1027, TITLE: Memory Aggregation Networks for Efficient Interactive Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Miao_Memory_Aggregation_Networks_for_Efficient_Interactive_Video_Object_Segmentation_CVPR_2020_paper.html
AUTHORS: Jiaxu Miao, Yunchao Wei, Yi Yang
HIGHLIGHT: In this work, we propose a unified framework, named Memory Aggregation Networks (MA-Net), to address the challenging iVOS in a more efficient way.

1028, TITLE: VQA With No Questions-Answers Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Vatashsky_VQA_With_No_Questions-Answers_Training_CVPR_2020_paper.html
AUTHORS: Ben-Zion Vatashsky, Shimon Ullman
HIGHLIGHT: We propose a novel method that consists of two main parts: generating a question graph representation, and an answering procedure, guided by the abstract structure of the question graph to invoke an extendable set of visual estimators.

1029, TITLE: Counting Out Time: Class Agnostic Video Repetition Counting in the Wild
http://openaccess.thecvf.com/content_CVPR_2020/html/Dwibedi_Counting_Out_Time_Class_Agnostic_Video_Repetition_Counting_in_the_CVPR_2020_paper.html
AUTHORS: Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, Andrew Zisserman

HIGHLIGHT:　　　　We present an approach for estimating the period with which an action is repeated in a video.

1030, TITLE:　　　　SaccadeNet: A Fast and Accurate Object Detector
http://openaccess.thecvf.com/content_CVPR_2020/html/Lan_SaccadeNet_A_Fast_and_Accurate_Object_Detector_CVPR_2020_paper.html
AUTHORS:　　　　Shiyi Lan, Zhou Ren, Yi Wu, Larry S. Davis, Gang Hua
HIGHLIGHT:　　　　In this paper, inspired by such mechanism, we propose a fast and accurate object detector called SaccadeNet.

1031, TITLE:　　　　Multi-Granularity Reference-Aided Attentive Feature Aggregation for Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Multi-Granularity_Reference-Aided_Attentive_Feature_Aggregation_for_Video-Based_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS:　　　　Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Zhibo Chen
HIGHLIGHT:　　　　In this paper, we propose an attentive feature aggregation module, namely Multi-Granularity Reference-aided Attentive Feature Aggregation (MG-RAFA), to delicately aggregate spatio-temporal features into a discriminative video-level feature representation.

1032, TITLE:　　　　Video Object Grounding Using Semantic Roles in Language Description
http://openaccess.thecvf.com/content_CVPR_2020/html/Sadhu_Video_Object_Grounding_Using_Semantic_Roles_in_Language_Description_CVPR_2020_paper.html
AUTHORS:　　　　Arka Sadhu, Kan Chen, Ram Nevatia
HIGHLIGHT:　　　　Here, we investigate the role of object relations in VOG and propose a novel framework VOGNet to encode multi-modal object relations via self-attention with relative position encoding.

1033, TITLE:　　　　Designing Network Design Spaces
http://openaccess.thecvf.com/content_CVPR_2020/html/Radosavovic_Designing_Network_Design_Spaces_CVPR_2020_paper.html
AUTHORS:　　　　Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, Piotr Dollar
HIGHLIGHT:　　　　In this work, we present a new network design paradigm.

1034, TITLE:　　　　12-in-1: Multi-Task Vision and Language Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_12-in-1_Multi-Task_Vision_and_Language_Representation_Learning_CVPR_2020_paper.html
AUTHORS:　　　　Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, Stefan Lee
HIGHLIGHT:　　　　In this work, we investigate these relationships between vision-and-language tasks by developing a large-scale, multi-task model.

1035, TITLE:　　　　MLCVNet: Multi-Level Context VoteNet for 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_MLCVNet_Multi-Level_Context_VoteNet_for_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS:　　　　Qian Xie, Yu-Kun Lai, Jing Wu, Zhoutao Wang, Yiming Zhang, Kai Xu, Jun Wang
HIGHLIGHT:　　　　In this paper, we address the 3D object detection task by capturing multi-level contextual information with the self-attention mechanism and multi-scale feature fusion.

1036, TITLE:　　　　Listen to Look: Action Recognition by Previewing Audio
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Listen_to_Look_Action_Recognition_by_Previewing_Audio_CVPR_2020_paper.html
AUTHORS:　　　　Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, Lorenzo Torresani
HIGHLIGHT:　　　　We propose a framework for efficient action recognition in untrimmed video that uses audio as a preview mechanism to eliminate both short-term and long-term visual redundancies.

1037, TITLE:　　　　Attention Convolutional Binary Neural Tree for Fine-Grained Visual Categorization
http://openaccess.thecvf.com/content_CVPR_2020/html/Ji_Attention_Convolutional_Binary_Neural_Tree_for_Fine-Grained_Visual_Categorization_CVPR_2020_paper.html
AUTHORS:　　　　Ruyi Ji, Longyin Wen, Libo Zhang, Dawei Du, Yanjun Wu, Chen Zhao, Xianglong Liu, Feiyue Huang
HIGHLIGHT:　　　　Specifically, we incorporate convolutional operations along edges of the tree structure, and use the routing functions in each node to determine the root-to-leaf computational paths within the tree.

1038, TITLE:　　　　Music Gesture for Visual Sound Separation
http://openaccess.thecvf.com/content_CVPR_2020/html/Gan_Music_Gesture_for_Visual_Sound_Separation_CVPR_2020_paper.html
AUTHORS:　　　　Chuang Gan, Deng Huang, Hang Zhao, Joshua B. Tenenbaum, Antonio Torralba

HIGHLIGHT: To address this, we propose "Music Gesture," a keypoint-based structured representation to explicitly model the body and finger movements of musicians when they perform music.

1039, TITLE: Referring Image Segmentation via Cross-Modal Progressive Comprehension
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Referring_Image_Segmentation_via_Cross-Modal_Progressive_Comprehension_CVPR_2020_paper.html
AUTHORS: Shaofei Huang, Tianrui Hui, Si Liu, Guanbin Li, Yunchao Wei, Jizhong Han, Luoqi Liu, Bo Li
HIGHLIGHT: In this paper, we propose a Cross-Modal Progressive Comprehension (CMPC) module and a Text-Guided Feature Exchange (TGFE) module to effectively address the challenging task.

1040, TITLE: Cloth in the Wind: A Case Study of Physical Measurement Through Simulation
http://openaccess.thecvf.com/content_CVPR_2020/html/Runia_Cloth_in_the_Wind_A_Case_Study_of_Physical_Measurement_CVPR_2020_paper.html
AUTHORS: Tom F. H. Runia, Kirill Gavrilyuk, Cees G. M. Snoek, Arnold W. M. Smeulders
HIGHLIGHT: In this paper, we propose to measure latent physical properties for cloth in the wind without ever having seen a real example before.

1041, TITLE: The Garden of Forking Paths: Towards Multi-Future Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Liang_The_Garden_of_Forking_Paths_Towards_Multi-Future_Trajectory_Prediction_CVPR_2020_paper.html
AUTHORS: Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, Alexander Hauptmann
HIGHLIGHT: This paper studies the problem of predicting the distribution over multiple possible future paths of people as they move through various visual scenes.

1042, TITLE: CentripetalNet: Pursuing High-Quality Keypoint Pairs for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_CentripetalNet_Pursuing_High-Quality_Keypoint_Pairs_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Zhiwei Dong, Guoxuan Li, Yue Liao, Fei Wang, Pengju Ren, Chen Qian
HIGHLIGHT: In this paper, we propose CentripetalNet which uses centripetal shift to pair corner keypoints from the same instance.

1043, TITLE: PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Shi_PV-RCNN_Point-Voxel_Feature_Set_Abstraction_for_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, Hongsheng Li
HIGHLIGHT: We present a novel and high-performance 3D object detection framework, named PointVoxel-RCNN (PV-RCNN), for accurate 3D object detection from point clouds.

1044, TITLE: Graph Embedded Pose Clustering for Anomaly Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Markovitz_Graph_Embedded_Pose_Clustering_for_Anomaly_Detection_CVPR_2020_paper.html
AUTHORS: Amir Markovitz, Gilad Sharir, Itamar Friedman, Lihi Zelnik-Manor, Shai Avidan
HIGHLIGHT: We propose a new method for anomaly detection of human actions.

1045, TITLE: Disp R-CNN: Stereo 3D Object Detection via Shape Prior Guided Instance Disparity Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Disp_R-CNN_Stereo_3D_Object_Detection_via_Shape_Prior_Guided_CVPR_2020_paper.html
AUTHORS: Jiaming Sun, Linghao Chen, Yiming Xie, Siyu Zhang, Qinhong Jiang, Xiaowei Zhou, Hujun Bao
HIGHLIGHT: In this paper, we propose a novel system named Disp R-CNN for 3D object detection from stereo images.

1046, TITLE: Deepstrip: High-Resolution Boundary Refinement
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Deepstrip_High-Resolution_Boundary_Refinement_CVPR_2020_paper.html
AUTHORS: Peng Zhou, Brian Price, Scott Cohen, Gregg Wilensky, Larry S. Davis
HIGHLIGHT: In this paper, we target refining the boundaries in high resolution images given low resolution masks.

1047, TITLE: Smoothing Adversarial Domain Attack and P-Memory Reconsolidation for Cross-Domain Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Smoothing_Adversarial_Domain_Attack_and_P-Memory_Reconsolidation_for_Cross-Domain_Person_CVPR_2020_paper.html
AUTHORS: Guangcong Wang, Jian-Huang Lai, Wenqi Liang, Guangrun Wang

HIGHLIGHT: To reduce the gap between the source and target domains, we propose a Smoothing Adversarial Domain Attack (SADA) approach that guides the source domain images to align the target domain images by using a trained camera classifier.

1048, TITLE: Meshed-Memory Transformer for Image Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Cornia_Meshed-Memory_Transformer_for_Image_Captioning_CVPR_2020_paper.html
AUTHORS: Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, Rita Cucchiara
HIGHLIGHT: With the aim of filling this gap, we present M2 - a Meshed Transformer with Memory for Image Captioning.

1049, TITLE: Learning From Noisy Anchors for One-Stage Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Learning_From_Noisy_Anchors_for_One-Stage_Object_Detection_CVPR_2020_paper.html
AUTHORS: Hengduo Li, Zuxuan Wu, Chen Zhu, Caiming Xiong, Richard Socher, Larry S. Davis
HIGHLIGHT: In this paper, we propose to mitigate noise incurred by imperfect label assignment such that the contributions of anchors are dynamically determined by a carefully constructed cleanliness score associated with each anchor.

1050, TITLE: Instance-Aware, Context-Focused, and Memory-Efficient Weakly Supervised Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Ren_Instance-Aware_Context-Focused_and_Memory-Efficient_Weakly_Supervised_Object_Detection_CVPR_2020_paper.html
AUTHORS: Zhongzheng Ren, Zhiding Yu, Xiaodong Yang, Ming-Yu Liu, Yong Jae Lee, Alexander G. Schwing, Jan Kautz
HIGHLIGHT: To target these issues we develop an instance-aware and context-focused unified framework.

1051, TITLE: Density-Based Clustering for 3D Object Detection in Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Ahmed_Density-Based_Clustering_for_3D_Object_Detection_in_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Syeda Mariam Ahmed, Chee Meng Chew
HIGHLIGHT: In this work, we introduce a novel approach for 3D object detection that is significant in two main aspects: a) cascaded modular approach that focuses the receptive field of each module on specific points in the point cloud, for improved feature learning and b) a class agnostic instance segmentation module that is initiated using unsupervised clustering.

1052, TITLE: Few-Shot Video Classification via Temporal Alignment
http://openaccess.thecvf.com/content_CVPR_2020/html/Cao_Few-Shot_Video_Classification_via_Temporal_Alignment_CVPR_2020_paper.html
AUTHORS: Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien-Yi Chang, Juan Carlos Niebles
HIGHLIGHT: In this paper, we propose the Ordered Temporal Alignment Module (OTAM), a novel few-shot learning framework that can learn to classify a previously unseen video.

1053, TITLE: Densely Connected Search Space for More Flexible Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Fang_Densely_Connected_Search_Space_for_More_Flexible_Neural_Architecture_Search_CVPR_2020_paper.html
AUTHORS: Jiemin Fang, Yuzhu Sun, Qian Zhang, Yuan Li, Wenyu Liu, Xinggang Wang
HIGHLIGHT: In this paper, we propose to search block counts and block widths by designing a densely connected search space, i.e., DenseNAS.

1054, TITLE: Fine-Grained Video-Text Retrieval With Hierarchical Graph Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Fine-Grained_Video-Text_Retrieval_With_Hierarchical_Graph_Reasoning_CVPR_2020_paper.html
AUTHORS: Shizhe Chen, Yida Zhao, Qin Jin, Qi Wu
HIGHLIGHT: To improve fine-grained video-text retrieval, we propose a Hierarchical Graph Reasoning (HGR) model, which decomposes video-text matching into global-to-local levels.

1055, TITLE: Warp to the Future: Joint Forecasting of Features and Feature Motion
http://openaccess.thecvf.com/content_CVPR_2020/html/Saric_Warp_to_the_Future_Joint_Forecasting_of_Features_and_Feature_CVPR_2020_paper.html
AUTHORS: Josip Saric, Marin Orsic, Tonci Antunovic, Sacha Vrazic, Sinisa Segvic
HIGHLIGHT: We propose to address this issue by complementing F2M forecasting with the classic F2F approach.

1056, TITLE: Network Adjustment: Channel Search Guided by FLOPs Utilization Ratio
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Network_Adjustment_Channel_Search_Guided_by_FLOPs_Utilization_Ratio_CVPR_2020_paper.html
AUTHORS: Zhengsu Chen, Jianwei Niu, Lingxi Xie, Xuefeng Liu, Longhui Wei, Qi Tian

HIGHLIGHT: This paper presents a new framework named network adjustment, which considers network accuracy as a function of FLOPs, so that under each network configuration, one can estimate the FLOPs utilization ratio (FUR) for each layer and use it to determine whether to increase or decrease the number of channels on the layer.

1057, TITLE: Where Does It Exist: Spatio-Temporal Video Grounding for Multi-Form Sentences
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Where_Does_It_Exist_Spatio-Temporal_Video_Grounding_for_Multi-Form_Sentences_CVPR_2020_paper.html
AUTHORS: Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, Lianli Gao
HIGHLIGHT: In this paper, we consider a novel task, Spatio-Temporal Video Grounding for Multi-Form Sentences (STVG).

1058, TITLE: Cross-Modal Cross-Domain Moment Alignment Network for Person Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Jing_Cross-Modal_Cross-Domain_Moment_Alignment_Network_for_Person_Search_CVPR_2020_paper.html
AUTHORS: Ya Jing, Wei Wang, Liang Wang, Tieniu Tan
HIGHLIGHT: Specially, we propose a moment alignment network (MAN) to solve the cross-modal cross-domain person search task in this paper.

1059, TITLE: Self-Training With Noisy Student Improves ImageNet Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_Self-Training_With_Noisy_Student_Improves_ImageNet_Classification_CVPR_2020_paper.html
AUTHORS: Qizhe Xie, Minh-Thang Luong, Eduard Hovy, Quoc V. Le
HIGHLIGHT: We present a simple self-training method that achieves 88.4% top-1 accuracy on ImageNet, which is 2.0% better than the state-of-the-art model that requires 3.5B weakly labeled Instagram images.

1060, TITLE: Learning Longterm Representations for Person Re-Identification Using Radio Signals
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_Learning_Longterm_Representations_for_Person_Re-Identification_Using_Radio_Signals_CVPR_2020_paper.html
AUTHORS: Lijie Fan, Tianhong Li, Rongyao Fang, Rumen Hristov, Yuan Yuan, Dina Katabi
HIGHLIGHT: In this paper, we introduce RF-ReID, a novel approach that harnesses radio frequency (RF) signals for longterm person ReID.

1061, TITLE: LatentFusion: End-to-End Differentiable Reconstruction and Rendering for Unseen Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Park_LatentFusion_End-to-End_Differentiable_Reconstruction_and_Rendering_for_Unseen_Object_Pose_CVPR_2020_paper.html
AUTHORS: Keunhong Park, Arsalan Mousavian, Yu Xiang, Dieter Fox
HIGHLIGHT: We propose a novel framework for 6D pose estimation of unseen objects.

1062, TITLE: Learning Instance Occlusion for Panoptic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lazarow_Learning_Instance_Occlusion_for_Panoptic_Segmentation_CVPR_2020_paper.html
AUTHORS: Justin Lazarow, Kwonjoon Lee, Kunyu Shi, Zhuowen Tu
HIGHLIGHT: To resolve this issue, we propose a branch that is tasked with modeling how two instance masks should overlap one another as a binary relation.

1063, TITLE: Vision-Dialog Navigation by Exploring Cross-Modal Memory
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Vision-Dialog_Navigation_by_Exploring_Cross-Modal_Memory_CVPR_2020_paper.html
AUTHORS: Yi Zhu, Fengda Zhu, Zhaohuan Zhan, Bingqian Lin, Jianbin Jiao, Xiaojun Chang, Xiaodan Liang
HIGHLIGHT: In this paper, we propose the Cross-modal Memory Network (CMN) for remembering and understanding the rich information relevant to historical navigation actions.

1064, TITLE: ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks
http://openaccess.thecvf.com/content_CVPR_2020/html/Shridhar_ALFRED_A_Benchmark_for_Interpreting_Grounded_Instructions_for_Everyday_Tasks_CVPR_2020_paper.html
AUTHORS: Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, Dieter Fox
HIGHLIGHT: We present ALFRED (Action Learning From Realistic Environments and Directives), a benchmark for learning a mapping from natural language instructions and egocentric vision to sequences of actions for household tasks.

1065, TITLE: NMS by Representative Region: Towards Crowded Pedestrian Detection by Proposal Pairing

http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_NMS_by_Representative_Region_Towards_Crowded_Pedestrian_De
tection_by_Proposal_CVPR_2020_paper.html
AUTHORS: Xin Huang, Zheng Ge, Zequn Jie, Osamu Yoshie
HIGHLIGHT: To avoid such a dilemma, this paper proposes a novel Representative Region NMS (R2NMS) approach
leveraging the less occluded visible parts, effectively removing the redundant boxes without bringing in many false positives.

1066, TITLE: Visual Commonsense R-CNN
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Visual_Commonsense_R-CNN_CVPR_2020_paper.html
AUTHORS: Tan Wang, Jianqiang Huang, Hanwang Zhang, Qianru Sun
HIGHLIGHT: We present a novel unsupervised feature representation learning method, Visual Commonsense Region-based
Convolutional Neural Network (VC R-CNN), to serve as an improved visual region encoder for high-level tasks such as captioning
and VQA.

1067, TITLE: What Deep CNNs Benefit From Global Covariance Pooling: An Optimization Perspective
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_What_Deep_CNNs_Benefit_From_Global_Covariance_Pooling_An_
Optimization_CVPR_2020_paper.html
AUTHORS: Qilong Wang, Li Zhang, Banggu Wu, Dongwei Ren, Peihua Li, Wangmeng Zuo, Qinghua Hu
HIGHLIGHT: In this paper, we make an attempt to understand what deep CNNs benefit from GCP in a viewpoint of
optimization.

1068, TITLE: EfficientDet: Scalable and Efficient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Tan_EfficientDet_Scalable_and_Efficient_Object_Detection_CVPR_2020_p
aper.html
AUTHORS: Mingxing Tan, Ruoming Pang, Quoc V. Le
HIGHLIGHT: In this paper, we systematically study neural network architecture design choices for object detection and
propose several key optimizations to improve efficiency.

1069, TITLE: Fast Template Matching and Update for Video Object Tracking and Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Fast_Template_Matching_and_Update_for_Video_Object_Tracking_an
d_CVPR_2020_paper.html
AUTHORS: Mingjie Sun, Jimin Xiao, Eng Gee Lim, Bingfeng Zhang, Yao Zhao
HIGHLIGHT: In this paper, the main task we aim to tackle is the multi-instance semi-supervised video object segmentation
across a sequence of frames where only the first-frame box-level ground-truth is provided.

1070, TITLE: Counterfactual Samples Synthesizing for Robust Visual Question Answering
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Counterfactual_Samples_Synthesizing_for_Robust_Visual_Question_
Answering_CVPR_2020_paper.html
AUTHORS: Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, Yueting Zhuang
HIGHLIGHT: To this end, we propose a model-agnostic Counterfactual Samples Synthesizing (CSS) training scheme.

1071, TITLE: Local-Global Video-Text Interactions for Temporal Grounding
http://openaccess.thecvf.com/content_CVPR_2020/html/Mun_Local-Global_Video-
Text_Interactions_for_Temporal_Grounding_CVPR_2020_paper.html
AUTHORS: Jonghwan Mun, Minsu Cho, Bohyung Han
HIGHLIGHT: We tackle this problem using a novel regression-based model that learns to extract a collection of mid-level
features for semantic phrases in a text query, which corresponds to important semantic entities described in the query (e.g., actors,
objects, and actions), and reflect bi-modal interactions between the linguistic features of the query and the visual features of the video
in multiple levels.

1072, TITLE: Set-Constrained Viterbi for Set-Supervised Action Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Set-Constrained_Viterbi_for_Set-
Supervised_Action_Segmentation_CVPR_2020_paper.html
AUTHORS: Jun Li, Sinisa Todorovic
HIGHLIGHT: Our first contribution is the formulation of a new set-constrained Viterbi algorithm (SCV).

1073, TITLE: Probabilistic Video Prediction From Noisy Data With a Posterior Confidence
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Probabilistic_Video_Prediction_From_Noisy_Data_With_a_Posterior
_Confidence_CVPR_2020_paper.html
AUTHORS: Yunbo Wang, Jiajun Wu, Mingsheng Long, Joshua B. Tenenbaum
HIGHLIGHT: In this paper, we propose to tackle this problem with an end-to-end trainable model named Bayesian Predictive
Network (BP-Net).

1074, TITLE: Beyond Short-Term Snippet: Video Relation Detection With Spatio-Temporal Global Context
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Beyond_Short-Term_Snippet_Video_Relation_Detection_With_Spatio-Temporal_Global_Context_CVPR_2020_paper.html
AUTHORS: Chenchen Liu, Yang Jin, Kehan Xu, Guoqiang Gong, Yadong Mu
HIGHLIGHT: To address these issues, this work proposes a novel sliding-window scheme to simultaneously predict short-term and long-term relationships.

1075, TITLE: Visual Grounding in Video for Unsupervised Word Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Sigurdsson_Visual_Grounding_in_Video_for_Unsupervised_Word_Translation_CVPR_2020_paper.html
AUTHORS: Gunnar A. Sigurdsson, Jean-Baptiste Alayrac, Aida Nematzadeh, Lucas Smaira, Mateusz Malinowski, Joao Carreira, Phil Blunsom, Andrew Zisserman
HIGHLIGHT: Our goal is to use visual grounding to improve unsupervised word mapping between languages.

1076, TITLE: Two Causal Principles for Improving Visual Dialog
http://openaccess.thecvf.com/content_CVPR_2020/html/Qi_Two_Causal_Principles_for_Improving_Visual_Dialog_CVPR_2020_paper.html
AUTHORS: Jiaxin Qi, Yulei Niu, Jianqiang Huang, Hanwang Zhang
HIGHLIGHT: This paper unravels the design tricks adopted by us, the champion team MReaL-BDAI, for Visual Dialog Challenge 2019: two causal principles for improving Visual Dialog (VisDial).

1077, TITLE: Spatio-Temporal Graph for Video Captioning With Knowledge Distillation
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Spatio-Temporal_Graph_for_Video_Captioning_With_Knowledge_Distillation_CVPR_2020_paper.html
AUTHORS: Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, Juan Carlos Niebles
HIGHLIGHT: In this paper, we propose a novel spatio-temporal graph model for video captioning that exploits object interactions in space and time.

1078, TITLE: A Real-Time Cross-Modality Correlation Filtering Method for Referring Expression Comprehension
http://openaccess.thecvf.com/content_CVPR_2020/html/Liao_A_Real-Time_Cross-Modality_Correlation_Filtering_Method_for_Referring_Expression_Comprehension_CVPR_2020_paper.html
AUTHORS: Yue Liao, Si Liu, Guanbin Li, Fei Wang, Yanjie Chen, Chen Qian, Bo Li
HIGHLIGHT: To this end, we propose a novel Realtime Cross-modality Correlation Filtering method (RCCF).

1079, TITLE: Better Captioning With Sequence-Level Exploration
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Better_Captioning_With_Sequence-Level_Exploration_CVPR_2020_paper.html
AUTHORS: Jia Chen, Qin Jin
HIGHLIGHT: In this work, we show the limitation of the current sequence-level learning objective for captioning tasks from both theory and empirical result.

1080, TITLE: Violin: A Large-Scale Dataset for Video-and-Language Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Violin_A_Large-Scale_Dataset_for_Video-and-Language_Inference_CVPR_2020_paper.html
AUTHORS: Jingzhou Liu, Wenhu Chen, Yu Cheng, Zhe Gan, Licheng Yu, Yiming Yang, Jingjing Liu
HIGHLIGHT: We introduce a new task, Video-and-Language Inference, for joint multimodal understanding of video and text.

1081, TITLE: RiFeGAN: Rich Feature Generation for Text-to-Image Synthesis From Prior Knowledge
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_RiFeGAN_Rich_Feature_Generation_for_Text-to-Image_Synthesis_From_Prior_Knowledge_CVPR_2020_paper.html
AUTHORS: Jun Cheng, Fuxiang Wu, Yanling Tian, Lei Wang, Dapeng Tao
HIGHLIGHT: To address this problem, we propose a novel rich feature generating text-to-image synthesis, called RiFeGAN, to enrich the given description.

1082, TITLE: Graph Structured Network for Image-Text Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Graph_Structured_Network_for_Image-Text_Matching_CVPR_2020_paper.html
AUTHORS: Chunxiao Liu, Zhendong Mao, Tianzhu Zhang, Hongtao Xie, Bin Wang, Yongdong Zhang
HIGHLIGHT: In this paper, we present a novel Graph Structured Matching Network (GSMN) to learn fine-grained correspondence.

1083, TITLE: Straight to the Point: Fast-Forwarding Videos via Reinforcement Learning Using Textual Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Ramos_Straight_to_the_Point_Fast-Forwarding_Videos_via_Reinforcement_Learning_Using_CVPR_2020_paper.html
AUTHORS: Washington Ramos, Michel Silva, Edson Araujo, Leandro Soriano Marcolino, Erickson Nascimento
HIGHLIGHT: In this paper, we present a novel methodology based on a reinforcement learning formulation to accelerate instructional videos.

1084, TITLE: Multi-Modality Cross Attention Network for Image and Sentence Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Multi-Modality_Cross_Attention_Network_for_Image_and_Sentence_Matching_CVPR_2020_paper.html
AUTHORS: Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, Feng Wu
HIGHLIGHT: Different from them, in this work, we propose a novel MultiModality Cross Attention (MMCA) Network for image and sentence matching by jointly modeling the intra-modality and inter-modality relationships of image regions and sentence words in a unified deep model.

1085, TITLE: Generalized ODIN: Detecting Out-of-Distribution Image Without Learning From Out-of-Distribution Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Hsu_Generalized_ODIN_Detecting_Out-of-Distribution_Image_Without_Learning_From_Out-of-Distribution_Data_CVPR_2020_paper.html
AUTHORS: Yen-Chang Hsu, Yilin Shen, Hongxia Jin, Zsolt Kira
HIGHLIGHT: We base our work on a popular method ODIN, proposing two strategies for freeing it from the needs of tuning with OoD data, while improving its OoD detection performance.

1086, TITLE: Learning Augmentation Network via Influence Functions
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_Learning_Augmentation_Network_via_Influence_Functions_CVPR_2020_paper.html
AUTHORS: Donghoon Lee, Hyunsin Park, Trung Pham, Chang D. Yoo
HIGHLIGHT: This paper considers an influence function that predicts how generalization performance, in terms of validation loss, is affected by a particular augmented training sample.

1087, TITLE: X-Linear Attention Networks for Image Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_X-Linear_Attention_Networks_for_Image_Captioning_CVPR_2020_paper.html
AUTHORS: Yingwei Pan, Ting Yao, Yehao Li, Tao Mei
HIGHLIGHT: In this paper, we introduce a unified attention block --- X-Linear attention block, that fully employs bilinear pooling to selectively capitalize on visual information or perform multi-modal reasoning.

1088, TITLE: Unsupervised Person Re-Identification via Multi-Label Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Unsupervised_Person_Re-Identification_via_Multi-Label_Classification_CVPR_2020_paper.html
AUTHORS: Dongkai Wang, Shiliang Zhang
HIGHLIGHT: This paper formulates unsupervised person ReID as a multi-label classification task to progressively seek true labels.

1089, TITLE: Overcoming Classifier Imbalance for Long-Tail Object Detection With Balanced Group Softmax
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Overcoming_Classifier_Imbalance_for_Long-Tail_Object_Detection_With_Balanced_Group_CVPR_2020_paper.html
AUTHORS: Yu Li, Tao Wang, Bingyi Kang, Sheng Tang, Chunfeng Wang, Jintao Li, Jiashi Feng
HIGHLIGHT: In this work, we propose a novel balanced group softmax (BAGS) module for balancing the classifiers within the detection frameworks through group-wise training.

1090, TITLE: What You See is What You Get: Exploiting Visibility for 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_What_You_See_is_What_You_Get_Exploiting_Visibility_for_CVPR_2020_paper.html
AUTHORS: Peiyun Hu, Jason Ziglar, David Held, Deva Ramanan
HIGHLIGHT: We argue that representing 2.5D data as collections of (x,y,z) points fundamentally destroys hidden information about freespace. In this paper, we demonstrate such knowledge can be efficiently recovered through 3D raycasting and readily incorporated into batch-based gradient learning.

1091, TITLE: Deep Structure-Revealed Network for Texture Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhai_Deep_Structure-Revealed_Network_for_Texture_Recognition_CVPR_2020_paper.html

AUTHORS: Wei Zhai, Yang Cao, Zheng-Jun Zha, HaiYong Xie, Feng Wu
HIGHLIGHT: To address this problem, we propose a novel Deep Structure-Revealed Network (DSR-Net) that leverages spatial dependency among the captured primitives as structural representation for texture recognition.

1092, TITLE: Online Knowledge Distillation via Collaborative Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Online_Knowledge_Distillation_via_Collaborative_Learning_CVPR_2020_paper.html
AUTHORS: Qiushan Guo, Xinjiang Wang, Yichao Wu, Zhipeng Yu, Ding Liang, Xiaolin Hu, Ping Luo
HIGHLIGHT: This work presents an efficient yet effective online Knowledge Distillation method via Collaborative Learning, termed KDCL, which is able to consistently improve the generalization ability of deep neural networks (DNNs) that have different learning capacities.

1093, TITLE: Dynamic Convolution: Attention Over Convolution Kernels
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Dynamic_Convolution_Attention_Over_Convolution_Kernels_CVPR_2020_paper.html
AUTHORS: Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, Zicheng Liu
HIGHLIGHT: To address this issue, we present Dynamic Convolution, a new design that increases model complexity without increasing the network depth or width.

1094, TITLE: 3DSSD: Point-Based 3D Single Stage Object Detector
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_3DSSD_Point-Based_3D_Single_Stage_Object_Detector_CVPR_2020_paper.html
AUTHORS: Zetong Yang, Yanan Sun, Shu Liu, Jiaya Jia
HIGHLIGHT: In this paper, we present a lightweight point-based 3D single stage object detector 3DSSD to achieve decent balance of accuracy and efficiency.

1095, TITLE: Deep Degradation Prior for Low-Quality Image Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Deep_Degradation_Prior_for_Low-Quality_Image_Classification_CVPR_2020_paper.html
AUTHORS: Yang Wang, Yang Cao, Zheng-Jun Zha, Jing Zhang, Zhiwei Xiong
HIGHLIGHT: To address this problem, this paper proposes a novel deep degradation prior for low-quality image classification.

1096, TITLE: ViBE: Dressing for Diverse Body Shapes
http://openaccess.thecvf.com/content_CVPR_2020/html/Hsiao_ViBE_Dressing_for_Diverse_Body_Shapes_CVPR_2020_paper.html
AUTHORS: Wei-Lin Hsiao, Kristen Grauman
HIGHLIGHT: We introduce ViBE, a VIsual Body-aware Embedding that captures clothing's affinity with different body shapes.

1097, TITLE: Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias
http://openaccess.thecvf.com/content_CVPR_2020/html/Singh_Dont_Judge_an_Object_by_Its_Context_Learning_to_Overcome_CVPR_2020_paper.html
AUTHORS: Krishna Kumar Singh, Dhruv Mahajan, Kristen Grauman, Yong Jae Lee, Matt Feiszli, Deepti Ghadiyaram
HIGHLIGHT: Our goal is to accurately recognize a category in the absence of its context, without compromising on performance when it co-occurs with context.

1098, TITLE: SESS: Self-Ensembling Semi-Supervised 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_SESS_Self-Ensembling_Semi-Supervised_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Na Zhao, Tat-Seng Chua, Gim Hee Lee
HIGHLIGHT: Inspired by the recent success of self-ensembling technique in semi-supervised image classification task, we propose SESS, a self-ensembling semi-supervised 3D object detection framework.

1099, TITLE: Combining Detection and Tracking for Human Pose Estimation in Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Combining_Detection_and_Tracking_for_Human_Pose_Estimation_in_Videos_CVPR_2020_paper.html
AUTHORS: Manchen Wang, Joseph Tighe, Davide Modolo
HIGHLIGHT: We propose a novel top-down approach that tackles the problem of multi-person human pose estimation and tracking in videos.

1100, TITLE: SAPIEN: A SimulAted Part-Based Interactive ENvironment

http://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_SAPIEN_A_SimulAted_Part-Based_Interactive_ENvironment_CVPR_2020_paper.html
AUTHORS: Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, Hao Su
HIGHLIGHT: SAPIEN enables various robotic vision and interaction tasks that require detailed part-level understanding.

1101, TITLE: RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_RandLA-Net_Efficient_Semantic_Segmentation_of_Large-Scale_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, Andrew Markham
HIGHLIGHT: In this paper, we introduce RandLA-Net, an efficient and lightweight neural architecture to directly infer per-point semantics for large-scale point clouds.

1102, TITLE: SurfelGAN: Synthesizing Realistic Sensor Data for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_SurfelGAN_Synthesizing_Realistic_Sensor_Data_for_Autonomous_Driving_CVPR_2020_paper.html
AUTHORS: Zhenpei Yang, Yuning Chai, Dragomir Anguelov, Yin Zhou, Pei Sun, Dumitru Erhan, Sean Rafferty, Henrik Kretzschmar
HIGHLIGHT: In this paper, we present a simple yet effective approach to generate realistic scenario sensor data, based only on a limited amount of lidar and camera data collected by an autonomous vehicle.

1103, TITLE: A Programmatic and Semantic Approach to Explaining and Debugging Neural Network Based Object Detectors
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_A_Programmatic_and_Semantic_Approach_to_Explaining_and_Debugging_Neural_CVPR_2020_paper.html
AUTHORS: Edward Kim, Divya Gopinath, Corina Pasareanu, Sanjit A. Seshia
HIGHLIGHT: In this paper, we present a programmatic and semantic approach to explaining, understanding, and debugging the correct and incorrect behaviors of a neural network based perception system.

1104, TITLE: Predicting Semantic Map Representations From Images Using Pyramid Occupancy Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Roddick_Predicting_Semantic_Map_Representations_From_Images_Using_Pyramid_Occupancy_Networks_CVPR_2020_paper.html
AUTHORS: Thomas Roddick, Roberto Cipolla
HIGHLIGHT: In this work we present a simple, unified approach for estimating these map representations directly from monocular images using a single end-to-end deep learning architecture.

1105, TITLE: Efficient Derivative Computation for Cumulative B-Splines on Lie Groups
http://openaccess.thecvf.com/content_CVPR_2020/html/Sommer_Efficient_Derivative_Computation_for_Cumulative_B-Splines_on_Lie_Groups_CVPR_2020_paper.html
AUTHORS: Christiane Sommer, Vladyslav Usenko, David Schubert, Nikolaus Demmel, Daniel Cremers
HIGHLIGHT: In this work we present an alternative derivation of time derivatives based on recurrence relations that needs $O(k)$ instead of $O(k^2)$ matrix operations (for a spline of order k) and results in simple and elegant expressions.

1106, TITLE: RL-CycleGAN: Reinforcement Learning Aware Simulation-to-Real
http://openaccess.thecvf.com/content_CVPR_2020/html/Rao_RL-CycleGAN_Reinforcement_Learning_Aware_Simulation-to-Real_CVPR_2020_paper.html
AUTHORS: Kanishka Rao, Chris Harris, Alex Irpan, Sergey Levine, Julian Ibarz, Mohi Khansari
HIGHLIGHT: In this paper, we introduce the RL-scene consistency loss for image translation, which ensures that the translation operation is invariant with respect to the Q-values associated with the image.

1107, TITLE: LiDARsim: Realistic LiDAR Simulation by Leveraging the Real World
http://openaccess.thecvf.com/content_CVPR_2020/html/Manivasagam_LiDARsim_Realistic_LiDAR_Simulation_by_Leveraging_the_Real_World_CVPR_2020_paper.html
AUTHORS: Sivabalan Manivasagam, Shenlong Wang, Kelvin Wong, Wenyuan Zeng, Mikita Sazanovich, Shuhan Tan, Bin Yang, Wei-Chiu Ma, Raquel Urtasun
HIGHLIGHT: We tackle the problem of producing realistic simulations of LiDAR point clouds, the sensor of preference for most self-driving vehicles.

1108, TITLE: Just Go With the Flow: Self-Supervised Scene Flow Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Mittal_Just_Go_With_the_Flow_Self-Supervised_Scene_Flow_Estimation_CVPR_2020_paper.html
AUTHORS: Himangi Mittal, Brian Okorn, David Held

HIGHLIGHT:　　　　As an alternative, we present a method of training scene flow that uses two self-supervised losses, based on nearest neighbors and cycle consistency.

1109, TITLE:　　　　TITAN: Future Forecast Using Action Priors

http://openaccess.thecvf.com/content_CVPR_2020/html/Malla_TITAN_Future_Forecast_Using_Action_Priors_CVPR_2020_paper.html

AUTHORS:　　　　Srikanth Malla, Behzad Dariush, Chiho Choi

HIGHLIGHT:　　　　In an attempt to address this problem, we introduce TITAN (Trajectory Inference using Targeted Action priors Network), a new model that incorporates prior positions, actions, and context to forecast future trajectory of agents and future ego-motion.

1110, TITLE:　　　　Robust Learning Through Cross-Task Consistency

http://openaccess.thecvf.com/content_CVPR_2020/html/Zamir_Robust_Learning_Through_Cross-Task_Consistency_CVPR_2020_paper.html

AUTHORS:　　　　Amir R. Zamir, Alexander Sax, Nikhil Cheerla, Rohan Suri, Zhangjie Cao, Jitendra Malik, Leonidas J. Guibas

HIGHLIGHT:　　　　We propose a flexible and fully computational framework for learning while enforcing Cross-Task Consistency (X-TAC).

1111, TITLE:　　　　Dynamic Refinement Network for Oriented and Densely Packed Object Detection

http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Dynamic_Refinement_Network_for_Oriented_and_Densely_Packed_Object_Detection_CVPR_2020_paper.html

AUTHORS:　　　　Xingjia Pan, Yuqiang Ren, Kekai Sheng, Weiming Dong, Haolei Yuan, Xiaowei Guo, Chongyang Ma, Changsheng Xu

HIGHLIGHT:　　　　To resolve the first two issues, we present a dynamic refinement network that consists of two novel components, i.e., a feature selection module (FSM) and a dynamic refinement head (DRH).

1112, TITLE:　　　　AOWS: Adaptive and Optimal Network Width Search With Latency Constraints

http://openaccess.thecvf.com/content_CVPR_2020/html/Berman_AOWS_Adaptive_and_Optimal_Network_Width_Search_With_Latency_Constraints_CVPR_2020_paper.html

AUTHORS:　　　　Maxim Berman, Leonid Pishchulin, Ning Xu, Matthew B. Blaschko, Gerard Medioni

HIGHLIGHT:　　　　We introduce a novel efficient one-shot NAS approach to optimally search for channel numbers, given latency constraints on a specific hardware.

1113, TITLE:　　　　High-Dimensional Convolutional Networks for Geometric Pattern Recognition

http://openaccess.thecvf.com/content_CVPR_2020/html/Choy_High-Dimensional_Convolutional_Networks_for_Geometric_Pattern_Recognition_CVPR_2020_paper.html

AUTHORS:　　　　Christopher Choy, Junha Lee, Rene Ranftl, Jaesik Park, Vladlen Koltun

HIGHLIGHT:　　　　In this work, we present high-dimensional convolutional networks for geometric pattern recognition problems that arise in 2D and 3D registration problems.

1114, TITLE:　　　　Filter Response Normalization Layer: Eliminating Batch Dependence in the Training of Deep Neural Networks

http://openaccess.thecvf.com/content_CVPR_2020/html/Singh_Filter_Response_Normalization_Layer_Eliminating_Batch_Dependence_in_the_Training_CVPR_2020_paper.html

AUTHORS:　　　　Saurabh Singh, Shankar Krishnan

HIGHLIGHT:　　　　In this paper we propose the Filter Response Normalization (FRN) layer, a novel combination of a normalization and an activation function, that can be used as a replacement for other normalizations and activations.

1115, TITLE:　　　　Deep Iterative Surface Normal Estimation

http://openaccess.thecvf.com/content_CVPR_2020/html/Lenssen_Deep_Iterative_Surface_Normal_Estimation_CVPR_2020_paper.html

AUTHORS:　　　　Jan Eric Lenssen, Christian Osendorfer, Jonathan Masci

HIGHLIGHT:　　　　This paper presents an end-to-end differentiable algorithm for robust and detail-preserving surface normal estimation on unstructured point-clouds.

1116, TITLE:　　　　Dataless Model Selection With the Deep Frame Potential

http://openaccess.thecvf.com/content_CVPR_2020/html/Murdock_Dataless_Model_Selection_With_the_Deep_Frame_Potential_CVPR_2020_paper.html

AUTHORS:　　　　Calvin Murdock, Simon Lucey

HIGHLIGHT:　　　　Building upon theoretical connections between deep learning and sparse approximation, we propose the deep frame potential: a measure of coherence that is approximately related to representation stability but has minimizers that depend only on network structure.

1117, TITLE: UNAS: Differentiable Architecture Search Meets Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Vahdat_UNAS_Differentiable_Architecture_Search_Meets_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS: Arash Vahdat, Arun Mallya, Ming-Yu Liu, Jan Kautz
HIGHLIGHT: In this work, we present UNAS, a unified framework for NAS, that encapsulates recent DNAS and RL-based approaches under one framework.

1118, TITLE: Local Context Normalization: Revisiting Local Normalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Ortiz_Local_Context_Normalization_Revisiting_Local_Normalization_CVPR_2020_paper.html
AUTHORS: Anthony Ortiz, Caleb Robinson, Dan Morris, Olac Fuentes, Christopher Kiekintveld, Md Mahmudulla Hassan, Nebojsa Jojic
HIGHLIGHT: We propose an algorithmic solution to make LCN efficient for arbitrary window sizes, even if every point in the image has a unique window.

1119, TITLE: ACNe: Attentive Context Normalization for Robust Permutation-Equivariant Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_ACNe_Attentive_Context_Normalization_for_Robust_Permutation-Equivariant_Learning_CVPR_2020_paper.html
AUTHORS: Weiwei Sun, Wei Jiang, Eduard Trulls, Andrea Tagliasacchi, Kwang Moo Yi
HIGHLIGHT: In this paper, we propose Attentive Context Normalization (ACN), a simple yet effective technique to build permutation-equivariant networks robust to outliers.

1120, TITLE: Learning Situational Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Ohn-Bar_Learning_Situational_Driving_CVPR_2020_paper.html
AUTHORS: Eshed Ohn-Bar, Aditya Prakash, Aseem Behl, Kashyap Chitta, Andreas Geiger
HIGHLIGHT: Our key idea is to learn a mixture model with a set of policies that can capture multiple driving modes.

1121, TITLE: From Depth What Can You See? Depth Completion via Auxiliary Image Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_From_Depth_What_Can_You_See_Depth_Completion_via_Auxiliary_CVPR_2020_paper.html
AUTHORS: Kaiyue Lu, Nick Barnes, Saeed Anwar, Liang Zheng
HIGHLIGHT: This paper continues this line of research and aims to overcome the above shortcomings.

1122, TITLE: Symmetry and Group in Attribute-Object Compositions
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Symmetry_and_Group_in_Attribute-Object_Compositions_CVPR_2020_paper.html
AUTHORS: Yong-Lu Li, Yue Xu, Xiaohan Mao, Cewu Lu
HIGHLIGHT: Incorporating the symmetry principle, a transformation framework inspired by group theory is built, i.e. SymNet.

1123, TITLE: Noise-Aware Fully Webly Supervised Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Noise-Aware_Fully_Webly_Supervised_Object_Detection_CVPR_2020_paper.html
AUTHORS: Yunhang Shen, Rongrong Ji, Zhiwei Chen, Xiaopeng Hong, Feng Zheng, Jianzhuang Liu, Mingliang Xu, Qi Tian
HIGHLIGHT: In this work, we propose an end-to-end framework to jointly learn webly supervised detectors and reduce the negative impact of noisy labels.

1124, TITLE: 3D Part Guided Image Editing for Fine-Grained Object Understanding
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_3D_Part_Guided_Image_Editing_for_Fine-Grained_Object_Understanding_CVPR_2020_paper.html
AUTHORS: Zongdai Liu, Feixiang Lu, Peng Wang, Hui Miao, Liangjun Zhang, Ruigang Yang, Bin Zhou
HIGHLIGHT: In this paper, we fill this important missing piece in autonomous driving by solving two critical issues.

1125, TITLE: STINet: Spatio-Temporal-Interactive Network for Pedestrian Detection and Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_STINet_Spatio-Temporal-Interactive_Network_for_Pedestrian_Detection_and_Trajectory_Prediction_CVPR_2020_paper.html
AUTHORS: Zhishuai Zhang, Jiyang Gao, Junhua Mao, Yukai Liu, Dragomir Anguelov, Congcong Li
HIGHLIGHT: In this work, we present a novel end-to-end two-stage network: Spatio-Temporal-Interactive Network (STINet).

1126, TITLE:		Rethinking Performance Estimation in Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Rethinking_Performance_Estimation_in_Neural_Architecture_Search_CVPR_2020_paper.html
AUTHORS:		Xiawu Zheng, Rongrong Ji, Qiang Wang, Qixiang Ye, Zhenguo Li, Yonghong Tian, Qi Tian
HIGHLIGHT:		In this paper, we provide a novel yet systematic rethinking of PE in a resource constrained regime, termed budgeted PE (BPE), which precisely and effectively estimates the performance of an architecture sampled from an architecture space.

1127, TITLE:		Feature-Metric Registration: A Fast Semi-Supervised Approach for Robust Point Cloud Registration Without Correspondences
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Feature-Metric_Registration_A_Fast_Semi-Supervised_Approach_for_Robust_Point_Cloud_CVPR_2020_paper.html
AUTHORS:		Xiaoshui Huang, Guofeng Mei, Jian Zhang
HIGHLIGHT:		We present a fast feature-metric point cloud registration framework, which enforces the optimisation of registration by minimising a feature-metric projection error without correspondences.

1128, TITLE:		Learning Multi-View Camera Relocalization With Graph Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Xue_Learning_Multi-View_Camera_Relocalization_With_Graph_Neural_Networks_CVPR_2020_paper.html
AUTHORS:		Fei Xue, Xin Wu, Shaojun Cai, Junqiu Wang
HIGHLIGHT:		We propose to construct a view graph to excavate the information of the whole given sequence for absolute camera pose estimation.

1129, TITLE:		MotionNet: Joint Perception and Motion Prediction for Autonomous Driving Based on Bird's Eye View Maps
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_MotionNet_Joint_Perception_and_Motion_Prediction_for_Autonomous_Driving_Based_CVPR_2020_paper.html
AUTHORS:		Pengxiang Wu, Siheng Chen, Dimitris N. Metaxas
HIGHLIGHT:		In this work, we propose an efficient deep model, called MotionNet, to jointly perform perception and motion prediction from 3D point clouds.

1130, TITLE:		EcoNAS: Finding Proxies for Economical Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_EcoNAS_Finding_Proxies_for_Economical_Neural_Architecture_Search_CVPR_2020_paper.html
AUTHORS:		Dongzhan Zhou, Xinchi Zhou, Wenwei Zhang, Chen Change Loy, Shuai Yi, Xuesen Zhang, Wanli Ouyang
HIGHLIGHT:		In this paper, we observe that most existing proxies exhibit different behaviors in maintaining the rank consistency among network candidates.

1131, TITLE:		Hit-Detector: Hierarchical Trinity Architecture Search for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Hit-Detector_Hierarchical_Trinity_Architecture_Search_for_Object_Detection_CVPR_2020_paper.html
AUTHORS:		Jianyuan Guo, Kai Han, Yunhe Wang, Chao Zhang, Zhaohui Yang, Han Wu, Xinghao Chen, Chang Xu
HIGHLIGHT:		To this end, we propose a hierarchical trinity search framework to simultaneously discover efficient architectures for all components (i.e. backbone, neck, and head) of object detector in an end-to-end manner.

1132, TITLE:		Geometrically Principled Connections in Graph Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Gong_Geometrically_Principled_Connections_in_Graph_Neural_Networks_CVPR_2020_paper.html
AUTHORS:		Shunwang Gong, Mehdi Bahri, Michael M. Bronstein, Stefanos Zafeiriou
HIGHLIGHT:		In this paper, we argue geometry should remain the primary driving force behind innovation in the emerging field of geometric deep learning.

1133, TITLE:		On Vocabulary Reliance in Scene Text Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Wan_On_Vocabulary_Reliance_in_Scene_Text_Recognition_CVPR_2020_paper.html
AUTHORS:		Zhaoyi Wan, Jielei Zhang, Liang Zhang, Jiebo Luo, Cong Yao
HIGHLIGHT:		In this paper, we establish an analytical framework, in which different datasets, metrics and module combinations for quantitative comparisons are devised, to conduct an in-depth study on the problem of vocabulary reliance in scene text recognition.

1134, TITLE:		Generating Accurate Pseudo-Labels in Semi-Supervised Learning and Avoiding Overconfident Predictions via Hermite Polynomial Activations
http://openaccess.thecvf.com/content_CVPR_2020/html/Lokhande_Generating_Accurate_Pseudo-Labels_in_Semi-Supervised_Learning_and_Avoiding_Overconfident_Predictions_CVPR_2020_paper.html

AUTHORS:	Vishnu Suresh Lokhande, Songwong Tasneeyapant, Abhay Venkatesh, Sathya N. Ravi, Vikas Singh
HIGHLIGHT:	Motivated by some of these results, we explore the use of Hermite polynomial expansions as a substitute for ReLUs in deep networks.

1135, TITLE:	GraspNet-1Billion: A Large-Scale Benchmark for General Object Grasping
http://openaccess.thecvf.com/content_CVPR_2020/html/Fang_GraspNet-1Billion_A_Large-Scale_Benchmark_for_General_Object_Grasping_CVPR_2020_paper.html
AUTHORS:	Hao-Shu Fang, Chenxi Wang, Minghao Gou, Cewu Lu
HIGHLIGHT:	In this work, we contribute a large-scale grasp pose detection dataset with a unified evaluation system.

1136, TITLE:	PFRL: Pose-Free Reinforcement Learning for 6D Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Shao_PFRL_Pose-Free_Reinforcement_Learning_for_6D_Pose_Estimation_CVPR_2020_paper.html
AUTHORS:	Jianzhun Shao, Yuhang Jiang, Gu Wang, Zhigang Li, Xiangyang Ji
HIGHLIGHT:	In this work, to get rid of the burden of 6D annotations, we formulate the 6D pose refinement as a Markov Decision Process and impose on the reinforcement learning approach with only 2D image annotations as weakly-supervised 6D pose information, via a delicate reward definition and a composite reinforced optimization method for efficient and effective policy training.

1137, TITLE:	Through Fog High-Resolution Imaging Using Millimeter Wave Radar
http://openaccess.thecvf.com/content_CVPR_2020/html/Guan_Through_Fog_High-Resolution_Imaging_Using_Millimeter_Wave_Radar_CVPR_2020_paper.html
AUTHORS:	Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, Haitham Hassanieh
HIGHLIGHT:	We introduce HawkEye, a system that leverages a cGAN architecture to recover high-frequency shapes from raw low-resolution mmWave heat-maps.

1138, TITLE:	Disentangling Physical Dynamics From Unknown Factors for Unsupervised Video Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Le_Guen_Disentangling_Physical_Dynamics_From_Unknown_Factors_for_Unsupervised_Video_Prediction_CVPR_2020_paper.html
AUTHORS:	Vincent Le Guen, Nicolas Thome
HIGHLIGHT:	Since physics is too restrictive for describing the full visual content of generic video sequences, we introduce PhyDNet, a two-branch deep architecture, which explicitly disentangles PDE dynamics from unknown complementary information.

1139, TITLE:	D2Det: Towards High Quality Object Detection and Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cao_D2Det_Towards_High_Quality_Object_Detection_and_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS:	Jiale Cao, Hisham Cholakkal, Rao Muhammad Anwer, Fahad Shahbaz Khan, Yanwei Pang, Ling Shao
HIGHLIGHT:	We propose a novel two-stage detection method, D2Det, that collectively addresses both precise localization and accurate classification.

1140, TITLE:	LiDAR-Based Online 3D Video Object Detection With Graph-Based Message Passing and Spatiotemporal Transformer Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Yin_LiDAR-Based_Online_3D_Video_Object_Detection_With_Graph-Based_Message_Passing_CVPR_2020_paper.html
AUTHORS:	Junbo Yin, Jianbing Shen, Chenye Guan, Dingfu Zhou, Ruigang Yang
HIGHLIGHT:	In this paper, we propose an end-to-end online 3D video object detector that operates on point cloud sequences.

1141, TITLE:	Orthogonal Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Orthogonal_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS:	Jiayun Wang, Yubei Chen, Rudrasis Chakraborty, Stella X. Yu
HIGHLIGHT:	We develop an efficient approach to impose filter orthogonality on a convolutional layer based on the doubly block-Toeplitz matrix representation of the convolutional kernel, instead of the common kernel orthogonality approach, which we show is only necessary but not sufficient for ensuring orthogonal convolutions.

1142, TITLE:	Self-Robust 3D Point Recognition via Gather-Vector Guidance
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Self-Robust_3D_Point_Recognition_via_Gather-Vector_Guidance_CVPR_2020_paper.html
AUTHORS:	Xiaoyi Dong, Dongdong Chen, Hang Zhou, Gang Hua, Weiming Zhang, Nenghai Yu
HIGHLIGHT:	In this paper, we look into the problem of 3D adversary attack, and propose to leverage the internal properties of the point clouds and the adversarial examples to design a new self-robust deep neural network (DNN) based 3D recognition systems.

1143, TITLE: VectorNet: Encoding HD Maps and Agent Dynamics From Vectorized Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_VectorNet_Encoding_HD_Maps_and_Agent_Dynamics_From_Vectorized_Representation_CVPR_2020_paper.html
AUTHORS: Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, Cordelia Schmid
HIGHLIGHT: This paper introduces VectorNet, a hierarchical graph neural network that first exploits the spatial locality of individual road components represented by vectors and then models the high-order interactions among all components.

1144, TITLE: ECA-Net: Efficient Channel Attention for Deep Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_ECA-Net_Efficient_Channel_Attention_for_Deep_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Qilong Wang, Banggu Wu, Pengfei Zhu, Peihua Li, Wangmeng Zuo, Qinghua Hu
HIGHLIGHT: To overcome the paradox of performance and complexity trade-off, this paper proposes an Efficient Channel Attention (ECA) module, which only involves a handful of parameters while bringing clear performance gain.

1145, TITLE: MTL-NAS: Task-Agnostic Neural Architecture Search Towards General-Purpose Multi-Task Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_MTL-NAS_Task-Agnostic_Neural_Architecture_Search_Towards_General-Purpose_Multi-Task_Learning_CVPR_2020_paper.html
AUTHORS: Yuan Gao, Haoping Bai, Zequn Jie, Jiayi Ma, Kui Jia, Wei Liu
HIGHLIGHT: We propose to incorporate neural architecture search (NAS) into general-purpose multi-task learning (GP-MTL).

1146, TITLE: PnPNet: End-to-End Perception and Prediction With Tracking in the Loop
http://openaccess.thecvf.com/content_CVPR_2020/html/Liang_PnPNet_End-to-End_Perception_and_Prediction_With_Tracking_in_the_Loop_CVPR_2020_paper.html
AUTHORS: Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, Raquel Urtasun
HIGHLIGHT: Towards this goal we propose PnPNet, an end-to-end model that takes as input sequential sensor data, and outputs at each time step object tracks and their future trajectories.

1147, TITLE: Revisiting the Sibling Head in Object Detector
http://openaccess.thecvf.com/content_CVPR_2020/html/Song_Revisiting_the_Sibling_Head_in_Object_Detector_CVPR_2020_paper.html
AUTHORS: Guanglu Song, Yu Liu, Xiaogang Wang
HIGHLIGHT: This paper provides the observation that the spatial misalignment between the two object functions in the sibling head can considerably hurt the training process, but this misalignment can be resolved by a very simple operator called task-aware spatial disentanglement (TSD).

1148, TITLE: Visual Reaction: Learning to Play Catch With Your Drone
http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_Visual_Reaction_Learning_to_Play_Catch_With_Your_Drone_CVPR_2020_paper.html
AUTHORS: Kuo-Hao Zeng, Roozbeh Mottaghi, Luca Weihs, Ali Farhadi
HIGHLIGHT: In this paper we address the problem of visual reaction: the task of interacting with dynamic environments where the changes in the environment are not necessarily caused by the agents itself.

1149, TITLE: Prime Sample Attention in Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Cao_Prime_Sample_Attention_in_Object_Detection_CVPR_2020_paper.html
AUTHORS: Yuhang Cao, Kai Chen, Chen Change Loy, Dahua Lin
HIGHLIGHT: In this work, we revisit this paradigm through a careful study on how different samples contribute to the overall performance measured in terms of mAP.

1150, TITLE: SpineNet: Learning Scale-Permuted Backbone for Recognition and Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Du_SpineNet_Learning_Scale-Permuted_Backbone_for_Recognition_and_Localization_CVPR_2020_paper.html
AUTHORS: Xianzhi Du, Tsung-Yi Lin, Pengchong Jin, Golnaz Ghiasi, Mingxing Tan, Yin Cui, Quoc V. Le, Xiaodan Song
HIGHLIGHT: In this paper, we argue encoder-decoder architecture is ineffective in generating strong multi-scale features because of the scale-decreased backbone.

1151, TITLE: KeyPose: Multi-View 3D Labeling and Keypoint Estimation for Transparent Objects
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_KeyPose_Multi-View_3D_Labeling_and_Keypoint_Estimation_for_Transparent_Objects_CVPR_2020_paper.html

AUTHORS: Xingyu Liu, Rico Jonschkowski, Anelia Angelova, Kurt Konolige
HIGHLIGHT: We address two problems: first, we establish an easy method for capturing and labeling 3D keypoints on desktop objects with an RGB camera; and second, we develop a deep neural network, called KeyPose, that learns to accurately predict object poses using 3D keypoints, from stereo input, and works even for transparent objects.

1152, TITLE: SegGCN: Efficient 3D Point Cloud Segmentation With Fuzzy Spherical Kernel
http://openaccess.thecvf.com/content_CVPR_2020/html/Lei_SegGCN_Efficient_3D_Point_Cloud_Segmentation_With_Fuzzy_Spherical_Kernel_CVPR_2020_paper.html
AUTHORS: Huan Lei, Naveed Akhtar, Ajmal Mian
HIGHLIGHT: Inspired by this observation, we incorporate a fuzzy mechanism into discrete convolutional kernels for 3D point clouds as our first major contribution.

1153, TITLE: nuScenes: A Multimodal Dataset for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Caesar_nuScenes_A_Multimodal_Dataset_for_Autonomous_Driving_CVPR_2020_paper.html
AUTHORS: Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, Oscar Beijbom
HIGHLIGHT: In this work we present nuTonomy scenes (nuScenes), the first dataset to carry the full autonomous vehicle sensor suite: 6 cameras, 5 radars and 1 lidar, all with full 360 degree field of view.

1154, TITLE: PVN3D: A Deep Point-Wise 3D Keypoints Voting Network for 6DoF Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/He_PVN3D_A_Deep_Point-Wise_3D_Keypoints_Voting_Network_for_6DoF_CVPR_2020_paper.html
AUTHORS: Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, Jian Sun
HIGHLIGHT: In this work, we present a novel data-driven method for robust 6DoF object pose estimation from a single RGBD image.

1155, TITLE: Probabilistic Pixel-Adaptive Refinement Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Wannenwetsch_Probabilistic_Pixel-Adaptive_Refinement_Networks_CVPR_2020_paper.html
AUTHORS: Anne S. Wannenwetsch, Stefan Roth
HIGHLIGHT: We introduce probabilistic pixel-adaptive convolutions (PPACs), which not only depend on image guidance data for filtering, but also respect the reliability of per-pixel predictions.

1156, TITLE: Discovering Human Interactions With Novel Objects via Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Discovering_Human_Interactions_With_Novel_Objects_via_Zero-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Suchen Wang, Kim-Hui Yap, Junsong Yuan, Yap-Peng Tan
HIGHLIGHT: We aim to detect human interactions with novel objects through zero-shot learning.

1157, TITLE: Equalization Loss for Long-Tailed Object Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Tan_Equalization_Loss_for_Long-Tailed_Object_Recognition_CVPR_2020_paper.html
AUTHORS: Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, Junjie Yan
HIGHLIGHT: In this work, we analyze this problem from a novel perspective: each positive sample of one category can be seen as a negative sample for other categories, making the tail categories receive more discouraging gradients.

1158, TITLE: Learning Depth-Guided Convolutions for Monocular 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Ding_Learning_Depth-Guided_Convolutions_for_Monocular_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Mingyu Ding, Yuqi Huo, Hongwei Yi, Zhe Wang, Jianping Shi, Zhiwu Lu, Ping Luo
HIGHLIGHT: In this work, instead of using pseudo-LiDAR representation, we improve the fundamental 2D fully convolutions by proposing a new local convolutional network (LCN), termed Depth-guided Dynamic-Depthwise-Dilated LCN (D4LCN), where the filters and their receptive fields can be automatically learned from image-based depth maps, making different pixels of different images have different filters.

1159, TITLE: Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather
http://openaccess.thecvf.com/content_CVPR_2020/html/Bijelic_Seeing_Through_Fog_Without_Seeing_Fog_Deep_Multimodal_Sensor_Fusion_CVPR_2020_paper.html
AUTHORS: Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, Felix Heide
HIGHLIGHT: To this end, we present a deep fusion network for robust fusion without a large corpus of labeled training data covering all asymmetric distortions.

1160, TITLE: Don't Even Look Once: Synthesizing Features for Zero-Shot Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Dont_Even_Look_Once_Synthesizing_Features_for_Zero-Shot_Detection_CVPR_2020_paper.html
AUTHORS: Pengkai Zhu, Hanxiao Wang, Venkatesh Saligrama
HIGHLIGHT: We propose a novel detection algorithm "Don't Even Look Once (DELO)," that synthesizes visual features for unseen objects and augments existing training algorithms to incorporate unseen object detection.

1161, TITLE: EPOS: Estimating 6D Pose of Objects With Symmetries
http://openaccess.thecvf.com/content_CVPR_2020/html/Hodan_EPOS_Estimating_6D_Pose_of_Objects_With_Symmetries_CVPR_2020_paper.html
AUTHORS: Tomas Hodan, Daniel Barath, Jiri Matas
HIGHLIGHT: We present a new method for estimating the 6D pose of rigid objects with available 3D models from a single RGB input image.

1162, TITLE: Train in Germany, Test in the USA: Making 3D Object Detectors Generalize
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Train_in_Germany_Test_in_the_USA_Making_3D_Object_CVPR_2020_paper.html
AUTHORS: Yan Wang, Xiangyu Chen, Yurong You, Li Erran Li, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger, Wei-Lun Chao
HIGHLIGHT: In this paper we consider the task of adapting 3D object detectors from one dataset to another.

1163, TITLE: Exploring Categorical Regularization for Domain Adaptive Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Exploring_Categorical_Regularization_for_Domain_Adaptive_Object_Detection_CVPR_2020_paper.html
AUTHORS: Chang-Dong Xu, Xing-Ran Zhao, Xin Jin, Xiu-Shen Wei
HIGHLIGHT: In this paper, we tackle the domain adaptive object detection problem, where the main challenge lies in significant domain gaps between source and target domains.

1164, TITLE: Neural Implicit Embedding for Point Cloud Analysis
http://openaccess.thecvf.com/content_CVPR_2020/html/Fujiwara_Neural_Implicit_Embedding_for_Point_Cloud_Analysis_CVPR_2020_paper.html
AUTHORS: Kent Fujiwara, Taiichi Hashimoto
HIGHLIGHT: We present a novel representation for point clouds that encapsulates the local characteristics of the underlying structure.

1165, TITLE: Pose-Guided Visible Part Matching for Occluded Person ReID
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Pose-Guided_Visible_Part_Matching_for_Occluded_Person_ReID_CVPR_2020_paper.html
AUTHORS: Shang Gao, Jingya Wang, Huchuan Lu, Zimo Liu
HIGHLIGHT: To address this issue, we propose a Pose-guided Visible Part Matching (PVPM) method that jointly learns the discriminative features with pose-guided attention and self-mines the part visibility in an end-to-end framework.

1166, TITLE: ContourNet: Taking a Further Step Toward Accurate Arbitrary-Shaped Scene Text Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_ContourNet_Taking_a_Further_Step_Toward_Accurate_Arbitrary-Shaped_Scene_Text_CVPR_2020_paper.html
AUTHORS: Yuxin Wang, Hongtao Xie, Zheng-Jun Zha, Mengting Xing, Zilong Fu, Yongdong Zhang
HIGHLIGHT: In this paper, we propose the ContourNet, which effectively handles these two problems taking a further step toward accurate arbitrary-shaped text detection.

1167, TITLE: Exploring Data Aggregation in Policy Learning for Vision-Based Urban Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Prakash_Exploring_Data_Aggregation_in_Policy_Learning_for_Vision-Based_Urban_Autonomous_CVPR_2020_paper.html
AUTHORS: Aditya Prakash, Aseem Behl, Eshed Ohn-Bar, Kashyap Chitta, Andreas Geiger
HIGHLIGHT: Our two key ideas are (1) to sample critical states from the collected on-policy data based on the utility they provide to the learned policy in terms of driving behavior, and (2) to incorporate a replay buffer which progressively focuses on the high uncertainty regions of the policy's state distribution.

1168, TITLE: Look-Into-Object: Self-Supervised Structure Modeling for Object Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Look-Into-Object_Self-Supervised_Structure_Modeling_for_Object_Recognition_CVPR_2020_paper.html

AUTHORS:	Mohan Zhou, Yalong Bai, Wei Zhang, Tiejun Zhao, Tao Mei
HIGHLIGHT:	In this paper, we propose to "look into object" (explicitly yet intrinsically model the object structure) through incorporating self-supervisions into the traditional framework.

1169, TITLE:	Recognizing Objects From Any View With Object and Viewer-Centered Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Recognizing_Objects_From_Any_View_With_Object_and_Viewer-Centered_Representations_CVPR_2020_paper.html
AUTHORS:	Sainan Liu, Vincent Nguyen, Isaac Rehg, Zhuowen Tu
HIGHLIGHT:	In this paper, we tackle an important task in computer vision: any view object recognition.

1170, TITLE:	Gated Channel Transformation for Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Gated_Channel_Transformation_for_Visual_Recognition_CVPR_2020_paper.html
AUTHORS:	Zongxin Yang, Linchao Zhu, Yu Wu, Yi Yang
HIGHLIGHT:	In this work, we propose a generally applicable transformation unit for visual recognition with deep convolutional neural networks.

1171, TITLE:	Non-Local Neural Networks With Grouped Bilinear Attentional Transforms
http://openaccess.thecvf.com/content_CVPR_2020/html/Chi_Non-Local_Neural_Networks_With_Grouped_Bilinear_Attentional_Transforms_CVPR_2020_paper.html
AUTHORS:	Lu Chi, Zehuan Yuan, Yadong Mu, Changhu Wang
HIGHLIGHT:	This work proposes a novel non-local operator. It is inspired by the attention mechanism of human visual system, which can quickly attend to important local parts in sight and suppress other less-relevant information.

1172, TITLE:	Generative-Discriminative Feature Representations for Open-Set Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Perera_Generative-Discriminative_Feature_Representations_for_Open-Set_Recognition_CVPR_2020_paper.html
AUTHORS:	Pramuditha Perera, Vlad I. Morariu, Rajiv Jain, Varun Manjunatha, Curtis Wigington, Vicente Ordonez, Vishal M. Patel
HIGHLIGHT:	We propose two techniques to force class activations of open-set samples to be low.

1173, TITLE:	RPM-Net: Robust Point Matching Using Learned Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Yew_RPM-Net_Robust_Point_Matching_Using_Learned_Features_CVPR_2020_paper.html
AUTHORS:	Zi Jian Yew, Gim Hee Lee
HIGHLIGHT:	In this paper, we propose the RPM-Net -- a less sensitive to initialization and more robust deep learning-based approach for rigid point cloud registration.

1174, TITLE:	Sideways: Depth-Parallel Training of Video Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Malinowski_Sideways_Depth-Parallel_Training_of_Video_Models_CVPR_2020_paper.html
AUTHORS:	Mateusz Malinowski, Grzegorz Swirszcz, Joao Carreira, Viorica Patraucean
HIGHLIGHT:	We propose Sideways, an approximate backpropagation scheme for training video models.

1175, TITLE:	Basis Prediction Networks for Effective Burst Denoising With Large Kernels
http://openaccess.thecvf.com/content_CVPR_2020/html/Xia_Basis_Prediction_Networks_for_Effective_Burst_Denoising_With_Large_Kernels_CVPR_2020_paper.html
AUTHORS:	Zhihao Xia, Federico Perazzi, Michael Gharbi, Kalyan Sunkavalli, Ayan Chakrabarti
HIGHLIGHT:	To this end, we introduce a novel basis prediction network that, given an input burst, predicts a set of global basis kernels --- shared within the image --- and the corresponding mixing coefficients --- which are specific to individual pixels.

1176, TITLE:	Private-kNN: Practical Differential Privacy for Computer Vision
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_Private-kNN_Practical_Differential_Privacy_for_Computer_Vision_CVPR_2020_paper.html
AUTHORS:	Yuqing Zhu, Xiang Yu, Manmohan Chandraker, Yu-Xiang Wang
HIGHLIGHT:	We propose a practically data-efficient scheme based on private release of k-nearest neighbor (kNN) queries, which altogether avoids splitting the training dataset.

1177, TITLE:	SP-NAS: Serial-to-Parallel Backbone Search for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_SP-NAS_Serial-to-Parallel_Backbone_Search_for_Object_Detection_CVPR_2020_paper.html

AUTHORS: Chenhan Jiang, Hang Xu, Wei Zhang, Xiaodan Liang, Zhenguo Li
HIGHLIGHT: In this paper, we propose a two-phase serial-to-parallel architecture search framework named SP-NAS towards a flexible task-oriented detection backbone.

1178, TITLE: Structure Aware Single-Stage 3D Object Detection From Point Cloud
http://openaccess.thecvf.com/content_CVPR_2020/html/He_Structure_Aware_Single-Stage_3D_Object_Detection_From_Point_Cloud_CVPR_2020_paper.html
AUTHORS: Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, Lei Zhang
HIGHLIGHT: In this work, we propose to improve the localization precision of single-stage detectors by explicitly leveraging the structure information of 3D point cloud.

1179, TITLE: "Looking at the Right Stuff" - Guided Semantic-Gaze for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Pal_Looking_at_the_Right_Stuff_-_Guided_Semantic-Gaze_for_Autonomous_CVPR_2020_paper.html
AUTHORS: Anwesan Pal, Sayan Mondal, Henrik I. Christensen
HIGHLIGHT: We propose a novel Semantics Augmented GazE (SAGE) detection approach that captures driving specific contextual information, in addition to the raw gaze.

1180, TITLE: What's Hidden in a Randomly Weighted Neural Network?
http://openaccess.thecvf.com/content_CVPR_2020/html/Ramanujan_Whats_Hidden_in_a_Randomly_Weighted_Neural_Network_CVPR_2020_paper.html
AUTHORS: Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, Mohammad Rastegari
HIGHLIGHT: We empirically show that as randomly weighted neural networks with fixed weights grow wider and deeper, an "untrained subnetwork" approaches a network with learned weights in accuracy.

1181, TITLE: Structured Multi-Hashing for Model Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Eban_Structured_Multi-Hashing_for_Model_Compression_CVPR_2020_paper.html
AUTHORS: Elad Eban, Yair Movshovitz-Attias, Hao Wu, Mark Sandler, Andrew Poon, Yerlan Idelbayev, Miguel A. Carreira-Perpinan
HIGHLIGHT: In this work we combine ideas from weight hashing and dimensionality reductions resulting in a simple and powerful structured multi-hashing method based on matrix products that allows direct control of model size of any deep network and is trained end-to-end.

1182, TITLE: DOPS: Learning to Detect 3D Objects and Predict Their 3D Shapes
http://openaccess.thecvf.com/content_CVPR_2020/html/Najibi_DOPS_Learning_to_Detect_3D_Objects_and_Predict_Their_3D_CVPR_2020_paper.html
AUTHORS: Mahyar Najibi, Guangda Lai, Abhijit Kundu, Zhichao Lu, Vivek Rathod, Thomas Funkhouser, Caroline Pantofaru, David Ross, Larry S. Davis, Alireza Fathi
HIGHLIGHT: We propose DOPS, a fast single-stage 3D object detection method for LIDAR data.

1183, TITLE: AutoTrack: Towards High-Performance Visual Tracking for UAV With Automatic Spatio-Temporal Regularization
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_AutoTrack_Towards_High-Performance_Visual_Tracking_for_UAV_With_Automatic_Spatio-Temporal_CVPR_2020_paper.html
AUTHORS: Yiming Li, Changhong Fu, Fangqiang Ding, Ziyuan Huang, Geng Lu
HIGHLIGHT: In this work, a novel approach is proposed to online automatically and adaptively learn spatio-temporal regularization term.

1184, TITLE: GP-NAS: Gaussian Process Based Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_GP-NAS_Gaussian_Process_Based_Neural_Architecture_Search_CVPR_2020_paper.html
AUTHORS: Zhihang Li, Teng Xi, Jiankang Deng, Gang Zhang, Shengzhao Wen, Ran He
HIGHLIGHT: In this paper, we aim to address three important questions in NAS: (1) How to measure the correlation between architectures and their performances? (2) How to evaluate the correlation between different architectures? (3) How to learn these correlations with a small number of samples?

1185, TITLE: NAS-FCOS: Fast Neural Architecture Search for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_NAS-FCOS_Fast_Neural_Architecture_Search_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Ning Wang, Yang Gao, Hao Chen, Peng Wang, Zhi Tian, Chunhua Shen, Yanning Zhang

HIGHLIGHT:	Here we propose to search for the decoder structure of object detectors with search efficiency being taken into consideration.

1186, TITLE:	TCTS: A Task-Consistent Two-Stage Framework for Person Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_TCTS_A_Task-Consistent_Two-Stage_Framework_for_Person_Search_CVPR_2020_paper.html
AUTHORS:	Cheng Wang, Bingpeng Ma, Hong Chang, Shiguang Shan, Xilin Chen
HIGHLIGHT:	To address the consistency problem, we introduce a Task-Consist Two-Stage (TCTS) person search framework, includes an identity-guided query (IDGQ) detector and a Detection Results Adapted (DRA) re-ID model.

1187, TITLE:	SCATTER: Selective Context Attentional Scene Text Recognizer
http://openaccess.thecvf.com/content_CVPR_2020/html/Litman_SCATTER_Selective_Context_Attentional_Scene_Text_Recognizer_CVPR_2020_paper.html
AUTHORS:	Ron Litman, Oron Anschel, Shahar Tsiper, Roee Litman, Shai Mazor, R. Manmatha
HIGHLIGHT:	In this paper, we introduce a novel architecture for STR, named Selective Context ATtentional Text Recognizer (SCATTER).

1188, TITLE:	Learning Canonical Shape Space for Category-Level 6D Object Pose and Size Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Learning_Canonical_Shape_Space_for_Category-Level_6D_Object_Pose_and_CVPR_2020_paper.html
AUTHORS:	Dengsheng Chen, Jun Li, Zheng Wang, Kai Xu
HIGHLIGHT:	We present a novel approach to category-level 6D object pose and size estimation.

1189, TITLE:	Hierarchical Scene Coordinate Classification and Regression for Visual Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Hierarchical_Scene_Coordinate_Classification_and_Regression_for_Visual_Localization_CVPR_2020_paper.html
AUTHORS:	Xiaotian Li, Shuzhe Wang, Yi Zhao, Jakob Verbeek, Juho Kannala
HIGHLIGHT:	In this work, we present a new hierarchical scene coordinate network to predict pixel scene coordinates in a coarse-to-fine manner from a single RGB image.

1190, TITLE:	MiLeNAS: Efficient Neural Architecture Search via Mixed-Level Reformulation
http://openaccess.thecvf.com/content_CVPR_2020/html/He_MiLeNAS_Efficient_Neural_Architecture_Search_via_Mixed-Level_Reformulation_CVPR_2020_paper.html
AUTHORS:	Chaoyang He, Haishan Ye, Li Shen, Tong Zhang
HIGHLIGHT:	To remedy this, this paper proposes MiLeNAS, a mixed-level reformulation for NAS that can be optimized efficiently and reliably.

1191, TITLE:	Scalable Uncertainty for Computer Vision With Functional Variational Inference
http://openaccess.thecvf.com/content_CVPR_2020/html/Carvalho_Scalable_Uncertainty_for_Computer_Vision_With_Functional_Variational_Inference_CVPR_2020_paper.html
AUTHORS:	Eduardo D. C. Carvalho, Ronald Clark, Andrea Nicastro, Paul H. J. Kelly
HIGHLIGHT:	By leveraging the structure of the induced covariance matrices, we propose numerically efficient algorithms which enable fast training in the context of high-dimensional tasks such as depth estimation and semantic segmentation.

1192, TITLE:	Uncertainty-Aware CNNs for Depth Completion: Uncertainty from Beginning to End
http://openaccess.thecvf.com/content_CVPR_2020/html/Eldesokey_Uncertainty-Aware_CNNs_for_Depth_Completion_Uncertainty_from_Beginning_to_End_CVPR_2020_paper.html
AUTHORS:	Abdelrahman Eldesokey, Michael Felsberg, Karl Holmquist, Michael Persson
HIGHLIGHT:	In this work, we thus focus on modeling the uncertainty of depth data in depth completion starting from the sparse noisy input all the way to the final prediction.

1193, TITLE:	Butterfly Transform: An Efficient FFT Based Neural Architecture Design
http://openaccess.thecvf.com/content_CVPR_2020/html/vahid_Butterfly_Transform_An_Efficient_FFT_Based_Neural_Architecture_Design_CVPR_2020_paper.html
AUTHORS:	Keivan Alizadeh vahid, Anish Prabhu, Ali Farhadi, Mohammad Rastegari
HIGHLIGHT:	In this paper, we show that extending the butterfly operations from the FFT algorithm to a general Butterfly Transform (BFT) can be beneficial in building an efficient block structure for CNN designs.

1194, TITLE:	A Certifiably Globally Optimal Solution to Generalized Essential Matrix Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_A_Certifiably_Globally_Optimal_Solution_to_Generalized_Essential_Matrix_Estimation_CVPR_2020_paper.html

AUTHORS: Ji Zhao, Wanting Xu, Laurent Kneip
HIGHLIGHT: We present a convex optimization approach for generalized essential matrix (GEM) estimation.

1195, TITLE: MUXConv: Information Multiplexing in Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_MUXConv_Information_Multiplexing_in_Convolutional_Neural_Networks_CVPR_2020_paper.html
AUTHORS: Zhichao Lu, Kalyanmoy Deb, Vishnu Naresh Boddeti
HIGHLIGHT: To overcome this limitation, we present MUXConv, a layer that is designed to increase the flow of information by progressively multiplexing channel and spatial information in the network, while mitigating computational complexity.

1196, TITLE: PointGMM: A Neural GMM Network for Point Clouds
http://openaccess.thecvf.com/content_CVPR_2020/html/Hertz_PointGMM_A_Neural_GMM_Network_for_Point_Clouds_CVPR_2020_paper.html
AUTHORS: Amir Hertz, Rana Hanocka, Raja Giryes, Daniel Cohen-Or
HIGHLIGHT: We present PointGMM, a neural network that learns to generate hGMMs which are characteristic of the shape class, and also coincide with the input point cloud.

1197, TITLE: Noisier2Noise: Learning to Denoise From Unpaired Noisy Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Moran_Noisier2Noise_Learning_to_Denoise_From_Unpaired_Noisy_Data_CVPR_2020_paper.html
AUTHORS: Nick Moran, Dan Schmidt, Yu Zhong, Patrick Coady
HIGHLIGHT: We present a method for training a neural network to perform image denoising without access to clean training examples or access to paired noisy training examples.

1198, TITLE: TRPLP - Trifocal Relative Pose From Lines at Points
http://openaccess.thecvf.com/content_CVPR_2020/html/Fabbri_TRPLP_-_Trifocal_Relative_Pose_From_Lines_at_Points_CVPR_2020_paper.html
AUTHORS: Ricardo Fabbri, Timothy Duff, Hongyi Fan, Margaret H. Regan, David da Costa de Pinho, Elias Tsigaridas, Charles W. Wampler, Jonathan D. Hauenstein, Peter J. Giblin, Benjamin Kimia, Anton Leykin, Tomas Pajdla
HIGHLIGHT: We present a method for solving two minimal problems for relative camera pose estimation from three views, which are based on three view correspondences of (i) three points and one line and (ii) three points and two lines through two of the points.

1199, TITLE: DSNAS: Direct Neural Architecture Search Without Parameter Retraining
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_DSNAS_Direct_Neural_Architecture_Search_Without_Parameter_Retraining_CVPR_2020_paper.html
AUTHORS: Shoukang Hu, Sirui Xie, Hehui Zheng, Chunxiao Liu, Jianping Shi, Xunying Liu, Dahua Lin
HIGHLIGHT: In this work, we propose a new problem definition for NAS, task-specific end-to-end, based on this observation.

1200, TITLE: MonoPair: Monocular 3D Object Detection Using Pairwise Spatial Relationships
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_MonoPair_Monocular_3D_Object_Detection_Using_Pairwise_Spatial_Relationships_CVPR_2020_paper.html
AUTHORS: Yongjian Chen, Lei Tai, Kai Sun, Mingyang Li
HIGHLIGHT: To this end, we propose a novel method to improve the monocular 3D object detection by considering the relationship of paired samples.

1201, TITLE: Regularization on Spatio-Temporally Smoothed Feature for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Regularization_on_Spatio-Temporally_Smoothed_Feature_for_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Jinhyung Kim, Seunghwan Cha, Dongyoon Wee, Soonmin Bae, Junmo Kim
HIGHLIGHT: In this paper, we propose Random Mean Scaling (RMS), a simple and effective regularization method, to relieve the overfitting problem in 3D residual networks.

1202, TITLE: Towards Accurate Scene Text Recognition With Semantic Reasoning Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Towards_Accurate_Scene_Text_Recognition_With_Semantic_Reasoning_Networks_CVPR_2020_paper.html
AUTHORS: Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, Errui Ding
HIGHLIGHT: To mitigate these limitations, we propose a novel end-to-end trainable framework named semantic reasoning network (SRN) for accurate scene text recognition, where a global semantic reasoning module (GSRM) is introduced to capture global semantic context through multi-way parallel transmission.

1203, TITLE:	Unsupervised Reinforcement Learning of Transferable Meta-Skills for Embodied Navigation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Unsupervised_Reinforcement_Learning_of_Transferable_Meta-Skills_for_Embodied_Navigation_CVPR_2020_paper.html
AUTHORS:	Juncheng Li, Xin Wang, Siliang Tang, Haizhou Shi, Fei Wu, Yueting Zhuang, William Yang Wang
HIGHLIGHT:	In this paper, we focus on visual navigation in the low-resource setting, where we have only a few training environments annotated with object information.

1204, TITLE:	Inferring Attention Shift Ranks of Objects for Image Saliency
http://openaccess.thecvf.com/content_CVPR_2020/html/Siris_Inferring_Attention_Shift_Ranks_of_Objects_for_Image_Saliency_CVPR_2020_paper.html
AUTHORS:	Avishek Siris, Jianbo Jiao, Gary K.L. Tam, Xianghua Xie, Rynson W.H. Lau
HIGHLIGHT:	Following psychological studies, in this paper, we propose to predict the saliency rank by inferring human attention shift.
Due to the lack of such data, we first construct a large-scale salient object ranking dataset.

1205, TITLE:	Camera On-Boarding for Person Re-Identification Using Hypothesis Transfer Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Ahmed_Camera_On-Boarding_for_Person_Re-Identification_Using_Hypothesis_Transfer_Learning_CVPR_2020_paper.html
AUTHORS:	Sk Miraj Ahmed, Aske R. Lejbolle, Rameswar Panda, Amit K. Roy-Chowdhury
HIGHLIGHT:	Rather, based on the fact that it is easy to store the learned re-identifications models, which mitigates any data privacy concern, we develop an efficient model adaptation approach using hypothesis transfer learning that aims to transfer the knowledge using only source models and limited labeled data, but without using any source camera data from the existing network.

1206, TITLE:	Joint Graph-Based Depth Refinement and Normal Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Rossi_Joint_Graph-Based_Depth_Refinement_and_Normal_Estimation_CVPR_2020_paper.html
AUTHORS:	Mattia Rossi, Mireille El Gheche, Andreas Kuhn, Pascal Frossard
HIGHLIGHT:	With these settings in mind, we devise a novel depth refinement framework that aims at recovering the underlying piece-wise planarity of those inverse depth maps associated to piece-wise planar scenes.

1207, TITLE:	DR Loss: Improving Object Detection by Distributional Ranking
http://openaccess.thecvf.com/content_CVPR_2020/html/Qian_DR_Loss_Improving_Object_Detection_by_Distributional_Ranking_CVPR_2020_paper.html
AUTHORS:	Qi Qian, Lei Chen, Hao Li, Rong Jin
HIGHLIGHT:	In this work, we propose a novel distributional ranking (DR) loss to handle the challenge.

1208, TITLE:	Self-Trained Deep Ordinal Regression for End-to-End Video Anomaly Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Pang_Self-Trained_Deep_Ordinal_Regression_for_End-to-End_Video_Anomaly_Detection_CVPR_2020_paper.html
AUTHORS:	Guansong Pang, Cheng Yan, Chunhua Shen, Anton van den Hengel, Xiao Bai
HIGHLIGHT:	By formulating a surrogate two-class ordinal regression task we devise an end-to-end trainable video anomaly detection approach that enables joint representation learning and anomaly scoring without manually labeled normal/abnormal data.

1209, TITLE:	Few-Shot Class-Incremental Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Tao_Few-Shot_Class-Incremental_Learning_CVPR_2020_paper.html
AUTHORS:	Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, Yihong Gong
HIGHLIGHT:	To address this problem, we represent the knowledge using a neural gas (NG) network, which can learn and preserve the topology of the feature manifold formed by different classes. On this basis, we propose the TOpology-Preserving knowledge InCrementer (TOPIC) framework.

1210, TITLE:	PolarMask: Single Shot Instance Segmentation With Polar Representation
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_PolarMask_Single_Shot_Instance_Segmentation_With_Polar_Representation_CVPR_2020_paper.html
AUTHORS:	Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, Ping Luo
HIGHLIGHT:	In this paper, we introduce an anchor-box free and single shot instance segmentation method, which is conceptually simple, fully convolutional and can be used by easily embedding it into most off-the-shelf detection methods.

1211, TITLE:	DeepEMD: Few-Shot Image Classification With Differentiable Earth Mover's Distance and Structured Classifiers
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_DeepEMD_Few-Shot_Image_Classification_With_Differentiable_Earth_Movers_Distance_and_CVPR_2020_paper.html
AUTHORS:	Chi Zhang, Yujun Cai, Guosheng Lin, Chunhua Shen

HIGHLIGHT: In this paper, we address the few-shot classification task from a new perspective of optimal matching between image regions.

1212, TITLE: Detection in Crowded Scenes: One Proposal, Multiple Predictions
http://openaccess.thecvf.com/content_CVPR_2020/html/Chu_Detection_in_Crowded_Scenes_One_Proposal_Multiple_Predictions_CVPR_2020_paper.html
AUTHORS: Xuangeng Chu, Anlin Zheng, Xiangyu Zhang, Jian Sun
HIGHLIGHT: We propose a simple yet effective proposal-based object detector, aiming at detecting highly-overlapped instances in crowded scenes.

1213, TITLE: Autolabeling 3D Objects With Differentiable Rendering of SDF Shape Priors
http://openaccess.thecvf.com/content_CVPR_2020/html/Zakharov_Autolabeling_3D_Objects_With_Differentiable_Rendering_of_SDF_Shape_Priors_CVPR_2020_paper.html
AUTHORS: Sergey Zakharov, Wadim Kehl, Arjun Bhargava, Adrien Gaidon
HIGHLIGHT: We present an automatic annotation pipeline to recover 9D cuboids and 3D shapes from pre-trained off-the-shelf 2D detectors and sparse LIDAR data.

1214, TITLE: Interactive Object Segmentation With Inside-Outside Guidance
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Interactive_Object_Segmentation_With_Inside-Outside_Guidance_CVPR_2020_paper.html
AUTHORS: Shiyin Zhang, Jun Hao Liew, Yunchao Wei, Shikui Wei, Yao Zhao
HIGHLIGHT: To achieve this, we propose an Inside-Outside Guidance (IOG) approach in this work.

1215, TITLE: Mnemonics Training: Multi-Class Incremental Learning Without Forgetting
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Mnemonics_Training_Multi-Class_Incremental_Learning_Without_Forgetting_CVPR_2020_paper.html
AUTHORS: Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, Qianru Sun
HIGHLIGHT: This paper proposes a novel and automatic framework we call mnemonics, where we parameterize exemplars and make them optimizable in an end-to-end manner.

1216, TITLE: Learning to Segment 3D Point Clouds in 2D Image Space
http://openaccess.thecvf.com/content_CVPR_2020/html/Lyu_Learning_to_Segment_3D_Point_Clouds_in_2D_Image_Space_CVPR_2020_paper.html
AUTHORS: Yecheng Lyu, Xinming Huang, Ziming Zhang
HIGHLIGHT: In contrast to the literature where local patterns in 3D point clouds are captured by customized convolutional operators, in this paper we study the problem of how to effectively and efficiently project such point clouds into a 2D image space so that traditional 2D convolutional neural networks (CNNs) such as U-Net can be applied for segmentation.

1217, TITLE: Smooth Shells: Multi-Scale Shape Registration With Functional Maps
http://openaccess.thecvf.com/content_CVPR_2020/html/Eisenberger_Smooth_Shells_Multi-Scale_Shape_Registration_With_Functional_Maps_CVPR_2020_paper.html
AUTHORS: Marvin Eisenberger, Zorah Lahner, Daniel Cremers
HIGHLIGHT: We propose a novel 3D shape correspondence method based on the iterative alignment of so-called smooth shells.

1218, TITLE: Self-Supervised Equivariant Attention Mechanism for Weakly Supervised Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Self-Supervised_Equivariant_Attention_Mechanism_for_Weakly_Supervised_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, Xilin Chen
HIGHLIGHT: In this paper, we propose a self-supervised equivariant attention mechanism (SEAM) to discover additional supervision and narrow the gap.

1219, TITLE: Efficient Neural Vision Systems Based on Convolutional Image Acquisition
http://openaccess.thecvf.com/content_CVPR_2020/html/Pad_Efficient_Neural_Vision_Systems_Based_on_Convolutional_Image_Acquisition_CVPR_2020_paper.html
AUTHORS: Pedram Pad, Simon Narduzzi, Clement Kundig, Engin Turetken, Siavash A. Bigdeli, L. Andrea Dunbar
HIGHLIGHT: In this paper, we tackle this fundamental challenge by introducing a hybrid optical-digital implementation of a convolutional neural network (CNN) based on engineering of the point spread function (PSF) of an optical imaging system.

1220, TITLE: Visual Chirality
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Visual_Chirality_CVPR_2020_paper.html

AUTHORS:		Zhiqiu Lin, Jin Sun, Abe Davis, Noah Snavely
HIGHLIGHT:		In this paper, we investigate how the statistics of visual data are changed by reflection.

1221, TITLE:		What Machines See Is Not What They Get: Fooling Scene Text Recognition Models With Adversarial Text Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_What_Machines_See_Is_Not_What_They_Get_Fooling_Scene_CVPR_2020_paper.html
AUTHORS:		Xing Xu, Jiefu Chen, Jinhui Xiao, Lianli Gao, Fumin Shen, Heng Tao Shen
HIGHLIGHT:		Specifically, we propose a novel and efficient optimization-based method that can be naturally integrated to different sequential prediction schemes, i.e., connectionist temporal classification (CTC) and attention mechanism.

1222, TITLE:		Dynamic Traffic Modeling From Overhead Imagery
http://openaccess.thecvf.com/content_CVPR_2020/html/Workman_Dynamic_Traffic_Modeling_From_Overhead_Imagery_CVPR_2020_paper.html
AUTHORS:		Scott Workman, Nathan Jacobs
HIGHLIGHT:		Instead, we propose an automatic approach for generating dynamic maps of traffic speeds using convolutional neural networks.

1223, TITLE:		Satellite Image Time Series Classification With Pixel-Set Encoders and Temporal Self-Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Garnot_Satellite_Image_Time_Series_Classification_With_Pixel-Set_Encoders_and_Temporal_CVPR_2020_paper.html
AUTHORS:		Vivien Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, Nesrine Chehata
HIGHLIGHT:		We propose an alternative approach in which the convolutional layers are advantageously replaced with encoders operating on unordered sets of pixels to exploit the typically coarse resolution of publicly available satellite images.

1224, TITLE:		DAVD-Net: Deep Audio-Aided Video Decompression of Talking Heads
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_DAVD-Net_Deep_Audio-Aided_Video_Decompression_of_Talking_Heads_CVPR_2020_paper.html
AUTHORS:		Xi Zhang, Xiaolin Wu, Xinliang Zhai, Xianye Ben, Chengjie Tu
HIGHLIGHT:		To address this problem, we present a novel deep convolutional neural network (DCNN) method for very low bit rate video reconstruction of talking heads.

1225, TITLE:		Learning When and Where to Zoom With Deep Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Uzkent_Learning_When_and_Where_to_Zoom_With_Deep_Reinforcement_Learning_CVPR_2020_paper.html
AUTHORS:		Burak Uzkent, Stefano Ermon
HIGHLIGHT:		In this direction, we propose PatchDrop a reinforcement learning approach to dynamically identify when and where to use/acquire high resolution data conditioned on the paired, cheap, low resolution images.

1226, TITLE:		Cross-Domain Detection via Graph-Induced Prototype Alignment
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Cross-Domain_Detection_via_Graph-Induced_Prototype_Alignment_CVPR_2020_paper.html
AUTHORS:		Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, Wenjun Zhang
HIGHLIGHT:		To mitigate these problems, we propose a Graph-induced Prototype Alignment (GPA) framework to seek for category-level domain alignment via elaborate prototype representations.

1227, TITLE:		Meta-Learning of Neural Architectures for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Elsken_Meta-Learning_of_Neural_Architectures_for_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS:		Thomas Elsken, Benedikt Staffler, Jan Hendrik Metzen, Frank Hutter
HIGHLIGHT:		To improve upon this, we propose MetaNAS, the first method which fully integrates NAS with gradient-based meta-learning.

1228, TITLE:		Towards Inheritable Models for Open-Set Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kundu_Towards_Inheritable_Models_for_Open-Set_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS:		Jogendra Nath Kundu, Naveen Venkat, Ambareesh Revanur, Rahul M V, R. Venkatesh Babu
HIGHLIGHT:		Addressing this, we introduce a practical DA paradigm where a source-trained model is used to facilitate adaptation in the absence of the source dataset in future.

1229, TITLE:		Learning From Synthetic Animals

http://openaccess.thecvf.com/content_CVPR_2020/html/Mu_Learning_From_Synthetic_Animals_CVPR_2020_paper.html
AUTHORS: Jiteng Mu, Weichao Qiu, Gregory D. Hager, Alan L. Yuille
HIGHLIGHT: In this paper, we use synthetic images and ground truth generated from CAD animal models to address this challenge.

1230, TITLE: Distilling Cross-Task Knowledge via Relationship Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_Distilling_Cross-Task_Knowledge_via_Relationship_Matching_CVPR_2020_paper.html
AUTHORS: Han-Jia Ye, Su Lu, De-Chuan Zhan
HIGHLIGHT: This paper deals with a general scenario reusing the knowledge from a cross-task teacher --- two models are targeting non-overlapping label spaces.

1231, TITLE: Open Compound Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Open_Compound_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Ziwei Liu, Zhongqi Miao, Xingang Pan, Xiaohang Zhan, Dahua Lin, Stella X. Yu, Boqing Gong
HIGHLIGHT: We propose a new approach based on two technical insights into OCDA: 1) a curriculum domain adaptation strategy to bootstrap generalization across domains in a data-driven self-organizing fashion and 2) a memory module to increase the model's agility towards novel domains.

1232, TITLE: Context Prior for Scene Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Context_Prior_for_Scene_Segmentation_CVPR_2020_paper.html
AUTHORS: Changqian Yu, Jingbo Wang, Changxin Gao, Gang Yu, Chunhua Shen, Nong Sang
HIGHLIGHT: In this work, we directly supervise the feature aggregation to distinguish the intra-class and interclass context clearly.

1233, TITLE: Tangent Images for Mitigating Spherical Distortion
http://openaccess.thecvf.com/content_CVPR_2020/html/Eder_Tangent_Images_for_Mitigating_Spherical_Distortion_CVPR_2020_paper.html
AUTHORS: Marc Eder, Mykhailo Shvets, John Lim, Jan-Michael Frahm
HIGHLIGHT: In this work, we propose "tangent images," a spherical image representation that facilitates transferable and scalable 360 degree computer vision.

1234, TITLE: Learning a Dynamic Map of Visual Appearance
http://openaccess.thecvf.com/content_CVPR_2020/html/Salem_Learning_a_Dynamic_Map_of_Visual_Appearance_CVPR_2020_paper.html
AUTHORS: Tawfiq Salem, Scott Workman, Nathan Jacobs
HIGHLIGHT: Every day billions of images capture this complex relationship, many of which are associated with precise time and location metadata. We propose to use these images to construct a global-scale, dynamic map of visual appearance attributes.

1235, TITLE: Webly Supervised Knowledge Embedding Model for Visual Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Webly_Supervised_Knowledge_Embedding_Model_for_Visual_Reasoning_CVPR_2020_paper.html
AUTHORS: Wenbo Zheng, Lan Yan, Chao Gou, Fei-Yue Wang
HIGHLIGHT: We present a two-stage approach for the task that can augment knowledge through an effective embedding model with weakly supervised web data.

1236, TITLE: Gradually Vanishing Bridge for Adversarial Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cui_Gradually_Vanishing_Bridge_for_Adversarial_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, Qi Tian
HIGHLIGHT: In this paper, we equip adversarial domain adaptation with Gradually Vanishing Bridge (GVB) mechanism on both generator and discriminator.

1237, TITLE: Active Speakers in Context
http://openaccess.thecvf.com/content_CVPR_2020/html/Alcazar_Active_Speakers_in_Context_CVPR_2020_paper.html
AUTHORS: Juan Leon Alcazar, Fabian Caba, Long Mai, Federico Perazzi, Joon-Young Lee, Pablo Arbelaez, Bernard Ghanem
HIGHLIGHT: This paper introduces the Active Speaker Context, a novel representation that models relationships between multiple speakers over long time horizons.

1238, TITLE: Panoptic-DeepLab: A Simple, Strong, and Fast Baseline for Bottom-Up Panoptic Segmentation

http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Panoptic-DeepLab_A_Simple_Strong_and_Fast_Baseline_for_Bottom-Up_Panoptic_CVPR_2020_paper.html
AUTHORS: Bowen Cheng, Maxwell D. Collins, Yukun Zhu, Ting Liu, Thomas S. Huang, Hartwig Adam, Liang-Chieh Chen
HIGHLIGHT: In this work, we introduce Panoptic-DeepLab, a simple, strong, and fast system for panoptic segmentation, aiming to establish a solid baseline for bottom-up methods that can achieve comparable performance of two-stage methods while yielding fast inference speed.

1239, TITLE: Inter-Region Affinity Distillation for Road Marking Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Hou_Inter-Region_Affinity_Distillation_for_Road_Marking_Segmentation_CVPR_2020_paper.html
AUTHORS: Yuenan Hou, Zheng Ma, Chunxiao Liu, Tak-Wai Hui, Chen Change Loy
HIGHLIGHT: In this work, we explore a novel knowledge distillation (KD) approach that can transfer 'knowledge' on scene structure more effectively from a teacher to a student model.

1240, TITLE: Unified Dynamic Convolutional Network for Super-Resolution With Variational Degradations
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Unified_Dynamic_Convolutional_Network_for_Super-Resolution_With_Variational_Degradations_CVPR_2020_paper.html
AUTHORS: Yu-Syuan Xu, Shou-Yao Roy Tseng, Yu Tseng, Hsien-Kai Kuo, Yi-Min Tsai
HIGHLIGHT: To fulfill this requirement, this paper proposes a unified network to accommodate the variations from inter-image (cross-image variations) and intra-image (spatial variations).

1241, TITLE: Making Better Mistakes: Leveraging Class Hierarchies With Deep Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Bertinetto_Making_Better_Mistakes_Leveraging_Class_Hierarchies_With_Deep_Networks_CVPR_2020_paper.html
AUTHORS: Luca Bertinetto, Romain Mueller, Konstantinos Tertikas, Sina Samangooei, Nicholas A. Lord
HIGHLIGHT: In this paper, we aim to renew interest in this problem by reviewing past approaches and proposing two simple methods which outperform the prior art under several metrics on two large datasets with complex class hierarchies: tieredImageNet and iNaturalist'19.

1242, TITLE: Data-Free Knowledge Amalgamation via Group-Stack Dual-GAN
http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_Data-Free_Knowledge_Amalgamation_via_Group-Stack_Dual-GAN_CVPR_2020_paper.html
AUTHORS: Jingwen Ye, Yixin Ji, Xinchao Wang, Xin Gao, Mingli Song
HIGHLIGHT: In this paper, we propose a data-free knowledge amalgamate strategy to craft a well-behaved multi-task student network from multiple single/multi-task teachers.

1243, TITLE: Screencast Tutorial Video Understanding
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Screencast_Tutorial_Video_Understanding_CVPR_2020_paper.html
AUTHORS: Kunpeng Li, Chen Fang, Zhaowen Wang, Seokhwan Kim, Hailin Jin, Yun Fu
HIGHLIGHT: In this paper, we propose visual understanding of screencast tutorials as a new research problem to the computer vision community.
We collect a new dataset of Adobe Photoshop video tutorials and annotate it with both low-level and high-level semantic labels.

1244, TITLE: DSGN: Deep Stereo Geometry Network for 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_DSGN_Deep_Stereo_Geometry_Network_for_3D_Object_Detection_CVPR_2020_paper.html
AUTHORS: Yilun Chen, Shu Liu, Xiaoyong Shen, Jiaya Jia
HIGHLIGHT: Our method, called Deep Stereo Geometry Network (DSGN), reduces this gap significantly by detecting 3D objects on a differentiable volumetric representation -- 3D geometric volume, which effectively encodes 3D geometric structure for 3D regular space.

1245, TITLE: Weakly-Supervised Salient Object Detection via Scribble Annotations
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Weakly-Supervised_Salient_Object_Detection_via_Scribble_Annotations_CVPR_2020_paper.html
AUTHORS: Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, Yuchao Dai
HIGHLIGHT: In this paper, we propose a weakly-supervised salient object detection model to learn saliency from such annotations.

1246, TITLE: Learning to Learn Single Domain Generalization
http://openaccess.thecvf.com/content_CVPR_2020/html/Qiao_Learning_to_Learn_Single_Domain_Generalization_CVPR_2020_paper.html

AUTHORS: Fengchun Qiao, Long Zhao, Xi Peng
HIGHLIGHT: We propose a new method named adversarial domain augmentation to solve this Out-of-Distribution (OOD) generalization problem.

1247, TITLE: Severity-Aware Semantic Segmentation With Reinforced Wasserstein Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Severity-Aware_Semantic_Segmentation_With_Reinforced_Wasserstein_Training_CVPR_2020_paper.html
AUTHORS: Xiaofeng Liu, Wenxuan Ji, Jane You, Georges El Fakhri, Jonghye Woo
HIGHLIGHT: To sidestep this, in this work, we propose to incorporate the severity-aware inter-class correlation into our Wasserstein training framework by configuring its ground distance matrix.

1248, TITLE: Boosting Few-Shot Learning With Adaptive Margin Loss
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Boosting_Few-Shot_Learning_With_Adaptive_Margin_Loss_CVPR_2020_paper.html
AUTHORS: Aoxue Li, Weiran Huang, Xu Lan, Jiashi Feng, Zhenguo Li, Liwei Wang
HIGHLIGHT: This paper proposes an adaptive margin principle to improve the generalization ability of metric-based meta-learning approaches for few-shot learning problems.

1249, TITLE: JA-POLS: A Moving-Camera Background Model via Joint Alignment and Partially-Overlapping Local Subspaces
http://openaccess.thecvf.com/content_CVPR_2020/html/Chelly_JA-POLS_A_Moving-Camera_Background_Model_via_Joint_Alignment_and_Partially-Overlapping_CVPR_2020_paper.html
AUTHORS: Irit Chelly, Vlad Winter, Dor Litvak, David Rosen, Oren Freifeld
HIGHLIGHT: Here we propose a purely-2D unsupervised modular method that systematically eliminates those issues.

1250, TITLE: AugFPN: Improving Multi-Scale Feature Learning for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_AugFPN_Improving_Multi-Scale_Feature_Learning_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Chaoxu Guo, Bin Fan, Qian Zhang, Shiming Xiang, Chunhong Pan
HIGHLIGHT: In this paper, we begin by first analyzing the design defects of feature pyramid in FPN, and then introduce a new feature pyramid architecture named AugFPN to address these problems.

1251, TITLE: xMUDA: Cross-Modal Unsupervised Domain Adaptation for 3D Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Jaritz_xMUDA_Cross-Modal_Unsupervised_Domain_Adaptation_for_3D_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Emilie Wirbel, Patrick Perez
HIGHLIGHT: In this work, we explore how to learn from multi-modality and propose cross-modal UDA (xMUDA) where we assume the presence of 2D images and 3D point clouds for 3D semantic segmentation.

1252, TITLE: Norm-Aware Embedding for Efficient Person Search
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Norm-Aware_Embedding_for_Efficient_Person_Search_CVPR_2020_paper.html
AUTHORS: Di Chen, Shanshan Zhang, Jian Yang, Bernt Schiele
HIGHLIGHT: To this end, We present a novel approach called Norm-Aware Embedding to disentangle the person embedding into norm and angle for detection and re-ID respectively, allowing for both effective and efficient multi-task training.

1253, TITLE: Intelligent Home 3D: Automatic 3D-House Design From Linguistic Descriptions Only
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Intelligent_Home_3D_Automatic_3D-House_Design_From_Linguistic_Descriptions_Only_CVPR_2020_paper.html
AUTHORS: Qi Chen, Qi Wu, Rui Tang, Yuhan Wang, Shuai Wang, Mingkui Tan
HIGHLIGHT: In this paper, we formulate it as a language conditioned visual content generation problem that is further divided into a floor plan generation and an interior texture (such as floor and wall) synthesis task.

1254, TITLE: Differential Treatment for Stuff and Things: A Simple Unsupervised Domain Adaptation Method for Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Differential_Treatment_for_Stuff_and_Things_A_Simple_Unsupervised_Domain_CVPR_2020_paper.html
AUTHORS: Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-mei Hwu, Thomas S. Huang, Honghui Shi
HIGHLIGHT: We consider the problem of unsupervised domain adaptation for semantic segmentation by easing the domain shift between the source domain (synthetic data) and the target domain (real data) in this work.

1255, TITLE: Robust Object Detection Under Occlusion With Context-Aware CompositionalNets
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Robust_Object_Detection_Under_Occlusion_With_Context-Aware_CompositionalNets_CVPR_2020_paper.html
AUTHORS: Angtian Wang, Yihong Sun, Adam Kortylewski, Alan L. Yuille
HIGHLIGHT: In this work, we propose to overcome two limitations of CompositionalNets which will enable them to detect partially occluded objects: 1) CompositionalNets, as well as other DCNN architectures, do not explicitly separate the representation of the context from the object itself.

1256, TITLE: IMRAM: Iterative Matching With Recurrent Attention Memory for Cross-Modal Image-Text Retrieval
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_IMRAM_Iterative_Matching_With_Recurrent_Attention_Memory_for_Cross-Modal_Image-Text_CVPR_2020_paper.html
AUTHORS: Hui Chen, Guiguang Ding, Xudong Liu, Zijia Lin, Ji Liu, Jungong Han
HIGHLIGHT: In this paper, to address such a deficiency, we propose an Iterative Matching with Recurrent Attention Memory (IMRAM) method, in which correspondences between images and texts are captured with multiple steps of alignments.

1257, TITLE: Domain-Aware Visual Bias Eliminating for Generalized Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Min_Domain-Aware_Visual_Bias_Eliminating_for_Generalized_Zero-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Shaobo Min, Hantao Yao, Hongtao Xie, Chaoqun Wang, Zheng-Jun Zha, Yongdong Zhang
HIGHLIGHT: In this paper, we propose a novel Domain-aware Visual Bias Eliminating (DVBE) network that constructs two complementary visual representations, i.e., semantic-free and semantic-aligned, to treat seen and unseen domains separately.

1258, TITLE: Semi-Supervised Semantic Segmentation With Cross-Consistency Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Ouali_Semi-Supervised_Semantic_Segmentation_With_Cross-Consistency_Training_CVPR_2020_paper.html
AUTHORS: Yassine Ouali, Celine Hudelot, Myriam Tami
HIGHLIGHT: In this paper, we present a novel cross-consistency based semi-supervised approach for semantic segmentation.

1259, TITLE: Learning to Learn Cropping Models for Different Aspect Ratio Requirements
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Learning_to_Learn_Cropping_Models_for_Different_Aspect_Ratio_Requirements_CVPR_2020_paper.html
AUTHORS: Debang Li, Junge Zhang, Kaiqi Huang
HIGHLIGHT: In this paper, we propose a meta-learning (learning to learn) based aspect ratio specified image cropping method called Mars, which can generate cropping results of different expected aspect ratios.

1260, TITLE: What Makes Training Multi-Modal Classification Networks Hard?
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_What_Makes_Training_Multi-Modal_Classification_Networks_Hard_CVPR_2020_paper.html
AUTHORS: Weiyao Wang, Du Tran, Matt Feiszli
HIGHLIGHT: This paper identifies two main causes for this performance drop: first, multi-modal networks are often prone to overfitting due to increased capacity. Second, different modalities overfit and generalize at different rates, so training them jointly with a single optimization strategy is sub-optimal. We address these two problems with a technique we call Gradient-Blending, which computes an optimal blending of modalities based on their overfitting behaviors.

1261, TITLE: Selective Transfer With Reinforced Transfer Network for Partial Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Selective_Transfer_With_Reinforced_Transfer_Network_for_Partial_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Zhihong Chen, Chao Chen, Zhaowei Cheng, Boyuan Jiang, Ke Fang, Xinyu Jin
HIGHLIGHT: In this paper, we propose a reinforced transfer network (RTNet), which utilizes both high-level and pixel-level information for PDA problem.

1262, TITLE: Semi-Supervised Semantic Image Segmentation With Self-Correcting Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Ibrahim_Semi-Supervised_Semantic_Image_Segmentation_With_Self-Correcting_Networks_CVPR_2020_paper.html
AUTHORS: Mostafa S. Ibrahim, Arash Vahdat, Mani Ranjbar, William G. Macready
HIGHLIGHT: In this paper, we introduce a principled semi-supervised framework that only use a small set of fully supervised images (having semantic segmentation labels and box labels) and a set of images with only object bounding box labels (we call it the weak-set).

1263, TITLE: Exemplar Normalization for Learning Deep Representation

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Exemplar_Normalization_for_Learning_Deep_Representation_CVPR_2020_paper.html
AUTHORS:		Ruimao Zhang, Zhanglin Peng, Lingyun Wu, Zhen Li, Ping Luo
HIGHLIGHT:		This work investigates a novel dynamic learning-to-normalize (L2N) problem by proposing Exemplar Normalization (EN), which is able to learn different normalization methods for different convolutional layers and image samples of a deep network.

1264, TITLE:		Imitative Non-Autoregressive Modeling for Trajectory Forecasting and Imputation
http://openaccess.thecvf.com/content_CVPR_2020/html/Qi_Imitative_Non-Autoregressive_Modeling_for_Trajectory_Forecasting_and_Imputation_CVPR_2020_paper.html
AUTHORS:		Mengshi Qi, Jie Qin, Yu Wu, Yi Yang
HIGHLIGHT:		To this end, we propose a novel imitative non-autoregressive modeling method to simultaneously handle the trajectory prediction task and the missing value imputation task.

1265, TITLE:		Multi-Modal Graph Neural Network for Joint Reasoning on Vision and Scene Text
http://openaccess.thecvf.com/content_CVPR_2020/html/Gao_Multi-Modal_Graph_Neural_Network_for_Joint_Reasoning_on_Vision_and_CVPR_2020_paper.html
AUTHORS:		Difei Gao, Ke Li, Ruiping Wang, Shiguang Shan, Xilin Chen
HIGHLIGHT:		Following this idea, we propose a novel VQA approach, Multi-Modal Graph Neural Network (MM-GNN).

1266, TITLE:		StereoGAN: Bridging Synthetic-to-Real Domain Gap by Joint Optimization of Domain Translation and Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_StereoGAN_Bridging_Synthetic-to-Real_Domain_Gap_by_Joint_Optimization_of_Domain_CVPR_2020_paper.html
AUTHORS:		Rui Liu, Chengxi Yang, Wenxiu Sun, Xiaogang Wang, Hongsheng Li
HIGHLIGHT:		In this paper, we propose an end-to-end training framework with domain translation and stereo matching networks to tackle this challenge.

1267, TITLE:		Self-Supervised Domain-Aware Generative Network for Generalized Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Self-Supervised_Domain-Aware_Generative_Network_for_Generalized_Zero-Shot_Learning_CVPR_2020_paper.html
AUTHORS:		Jiamin Wu, Tianzhu Zhang, Zheng-Jun Zha, Jiebo Luo, Yongdong Zhang, Feng Wu
HIGHLIGHT:		To address this issue, we propose an end-to-end Self-supervised Domain-aware Generative Network (SDGN) by integrating self-supervised learning into feature generating model for unbiased GZSL.

1268, TITLE:		Sparse Layered Graphs for Multi-Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Jeppesen_Sparse_Layered_Graphs_for_Multi-Object_Segmentation_CVPR_2020_paper.html
AUTHORS:		Niels Jeppesen, Anders N. Christensen, Vedrana A. Dahl, Anders B. Dahl
HIGHLIGHT:		We introduce the novel concept of a Sparse Layered Graph (SLG) for s-t graph cut segmentation of image data.

1269, TITLE:		Visual-Semantic Matching by Exploring High-Order Attention and Distraction
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Visual-Semantic_Matching_by_Exploring_High-Order_Attention_and_Distraction_CVPR_2020_paper.html
AUTHORS:		Yongzhi Li, Duo Zhang, Yadong Mu
HIGHLIGHT:		In this work, we address this task from two previously-ignored aspects: high-order semantic information (e.g., object-predicate-subject triplet, object-attribute pair) and visual distraction (i.e., despite the high relevance to textual query, images may also contain many prominent distracting objects or visual relations).

1270, TITLE:		End-to-End 3D Point Cloud Instance Segmentation Without Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Jiang_End-to-End_3D_Point_Cloud_Instance_Segmentation_Without_Detection_CVPR_2020_paper.html
AUTHORS:		Haiyong Jiang, Feilong Yan, Jianfei Cai, Jianmin Zheng, Jun Xiao
HIGHLIGHT:		In this paper, we introduce a novel framework to enable end-to-end instance segmentation without detection and a separate step of grouping.

1271, TITLE:		Deep Adversarial Decomposition: A Unified Framework for Separating Superimposed Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Zou_Deep_Adversarial_Decomposition_A_Unified_Framework_for_Separating_Superimposed_Images_CVPR_2020_paper.html
AUTHORS:		Zhengxia Zou, Sen Lei, Tianyang Shi, Zhenwei Shi, Jieping Ye
HIGHLIGHT:		We propose a unified framework named "deep adversarial decomposition" for single superimposed image separation.

1272, TITLE: Differentiable Adaptive Computation Time for Visual Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Eyzaguirre_Differentiable_Adaptive_Computation_Time_for_Visual_Reasoning_CVPR_2020_paper.html
AUTHORS: Cristobal Eyzaguirre, Alvaro Soto
HIGHLIGHT: This paper presents a novel attention-based algorithm for achieving adaptive computation called DACT, which, unlike existing ones, is end-to-end differentiable.

1273, TITLE: DeepLPF: Deep Local Parametric Filters for Image Enhancement
http://openaccess.thecvf.com/content_CVPR_2020/html/Moran_DeepLPF_Deep_Local_Parametric_Filters_for_Image_Enhancement_CVPR_2020_paper.html
AUTHORS: Sean Moran, Pierre Marza, Steven McDonagh, Sarah Parisot, Gregory Slabaugh
HIGHLIGHT: In this paper, we introduce a novel approach to automatically enhance images using learned spatially local filters of three different types (Elliptical Filter, Graduated Filter, Polynomial Filter).

1274, TITLE: Instance Credibility Inference for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Instance_Credibility_Inference_for_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Yikai Wang, Chengming Xu, Chen Liu, Li Zhang, Yanwei Fu
HIGHLIGHT: In contrast, this paper presents a simple statistical approach, dubbed Instance Credibility Inference (ICI) to exploit the distribution support of unlabeled instances for few-shot learning.

1275, TITLE: Learning From Web Data With Self-Organizing Memory Module
http://openaccess.thecvf.com/content_CVPR_2020/html/Tu_Learning_From_Web_Data_With_Self-Organizing_Memory_Module_CVPR_2020_paper.html
AUTHORS: Yi Tu, Li Niu, Junjie Chen, Dawei Cheng, Liqing Zhang
HIGHLIGHT: In this paper, we propose a novel method, which is capable of handling these two types of noises together, without the supervision of clean images in the training stage.

1276, TITLE: TransMatch: A Transfer-Learning Scheme for Semi-Supervised Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_TransMatch_A_Transfer-Learning_Scheme_for_Semi-Supervised_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Zhongjie Yu, Lin Chen, Zhongwei Cheng, Jiebo Luo
HIGHLIGHT: In this paper, we propose a new transfer-learning framework for semi-supervised few-shot learning to fully utilize the auxiliary information from labeled base-class data and unlabeled novel-class data.

1277, TITLE: Learning the Redundancy-Free Features for Generalized Zero-Shot Object Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_Learning_the_Redundancy-Free_Features_for_Generalized_Zero-Shot_Object_Recognition_CVPR_2020_paper.html
AUTHORS: Zongyan Han, Zhenyong Fu, Jian Yang
HIGHLIGHT: To reduce the superfluous information in the fine-grained objects, in this paper, we propose to learn the redundancy-free features for generalized zero-shot learning.

1278, TITLE: Neural Topological SLAM for Visual Navigation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chaplot_Neural_Topological_SLAM_for_Visual_Navigation_CVPR_2020_paper.html
AUTHORS: Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, Saurabh Gupta
HIGHLIGHT: This paper studies the problem of image-goal navigation which involves navigating to the location indicated by a goal image in a novel previously unseen environment.

1279, TITLE: WaveletStereo: Learning Wavelet Coefficients of Disparity Map in Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_WaveletStereo_Learning_Wavelet_Coefficients_of_Disparity_Map_in_Stereo_Matching_CVPR_2020_paper.html
AUTHORS: Menglong Yang, Fangrui Wu, Wei Li
HIGHLIGHT: This paper proposes a novel stereo matching method called WaveletStereo, which learns the wavelet coefficients of the disparity rather than the disparity itself.

1280, TITLE: Robust Superpixel-Guided Attentional Adversarial Attack
http://openaccess.thecvf.com/content_CVPR_2020/html/Dong_Robust_Superpixel-Guided_Attentional_Adversarial_Attack_CVPR_2020_paper.html

AUTHORS: Xiaoyi Dong, Jiangfan Han, Dongdong Chen, Jiayang Liu, Huanyu Bian, Zehua Ma, Hongsheng Li, Xiaogang Wang, Weiming Zhang, Nenghai Yu
HIGHLIGHT: Based on these two considerations, we propose the first robust superpixel-guided attentional adversarial attack method.

1281, TITLE: BEDSR-Net: A Deep Shadow Removal Network From a Single Document Image
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_BEDSR-Net_A_Deep_Shadow_Removal_Network_From_a_Single_Document_CVPR_2020_paper.html
AUTHORS: Yun-Hsuan Lin, Wen-Chin Chen, Yung-Yu Chuang
HIGHLIGHT: This paper proposes the Background Estimation Document Shadow Removal Network (BEDSR-Net), the first deep network specifically designed for document image shadow removal.

1282, TITLE: Cross-Domain Document Object Detection: Benchmark Suite and Method
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Cross-Domain_Document_Object_Detection_Benchmark_Suite_and_Method_CVPR_2020_paper.html
AUTHORS: Kai Li, Curtis Wigington, Chris Tensmeyer, Handong Zhao, Nikolaos Barmpalios, Vlad I. Morariu, Varun Manjunatha, Tong Sun, Yun Fu
HIGHLIGHT: We investigate cross-domain DOD, where the goal is to learn a detector for the target domain using labeled data from the source domain and only unlabeled data from the target domain.

1283, TITLE: Explaining Knowledge Distillation by Quantifying the Knowledge
http://openaccess.thecvf.com/content_CVPR_2020/html/Cheng_Explaining_Knowledge_Distillation_by_Quantifying_the_Knowledge_CVPR_2020_paper.html
AUTHORS: Xu Cheng, Zhefan Rao, Yilan Chen, Quanshi Zhang
HIGHLIGHT: This paper presents a method to interpret the success of knowledge distillation by quantifying and analyzing task-relevant and task-irrelevant visual concepts that are encoded in intermediate layers of a deep neural network (DNN).

1284, TITLE: Exploring Bottom-Up and Top-Down Cues With Attentive Learning for Webly Supervised Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Exploring_Bottom-Up_and_Top-Down_Cues_With_Attentive_Learning_for_Webly_CVPR_2020_paper.html
AUTHORS: Zhonghua Wu, Qingyi Tao, Guosheng Lin, Jianfei Cai
HIGHLIGHT: Within our approach, we introduce a bottom-up mechanism based on the well-trained fully supervised object detector (i.e. Faster RCNN) as an object region estimator for web images by recognizing the common objectiveness shared by base and novel classes.

1285, TITLE: Enhancing Generic Segmentation With Learned Region Representations
http://openaccess.thecvf.com/content_CVPR_2020/html/Isaacs_Enhancing_Generic_Segmentation_With_Learned_Region_Representations_CVPR_2020_paper.html
AUTHORS: Or Isaacs, Oran Shayer, Michael Lindenbaum
HIGHLIGHT: We propose an alternative approach called Deep Generic Segmentation (DGS) and try to follow the path used for semantic segmentation.

1286, TITLE: Adaptive Hierarchical Down-Sampling for Point Cloud Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Nezhadarya_Adaptive_Hierarchical_Down-Sampling_for_Point_Cloud_Classification_CVPR_2020_paper.html
AUTHORS: Ehsan Nezhadarya, Ehsan Taghavi, Ryan Razani, Bingbing Liu, Jun Luo
HIGHLIGHT: In this paper, we propose a novel deterministic, adaptive, permutation-invariant down-sampling layer, called Critical Points Layer (CPL), which learns to reduce the number of points in an unordered point cloud while retaining the important (critical) ones.

1287, TITLE: FBNetV2: Differentiable Neural Architecture Search for Spatial and Channel Dimensions
http://openaccess.thecvf.com/content_CVPR_2020/html/Wan_FBNetV2_Differentiable_Neural_Architecture_Search_for_Spatial_and_Channel_Dimensions_CVPR_2020_paper.html
AUTHORS: Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, Joseph E. Gonzalez
HIGHLIGHT: To address this bottleneck, we propose a memory and computationally efficient DNAS variant: DMaskingNAS.

1288, TITLE: Learning Texture Invariant Representation for Domain Adaptation of Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Learning_Texture_Invariant_Representation_for_Domain_Adaptation_of_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Myeongjin Kim, Hyeran Byun

HIGHLIGHT: In this paper, considering the fundamental difference between the two domains as the texture, we propose a method to adapt to the target domain's texture.

1289, TITLE: Putting Visual Object Recognition in Context
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Putting_Visual_Object_Recognition_in_Context_CVPR_2020_paper.html
AUTHORS: Mengmi Zhang, Claire Tseng, Gabriel Kreiman
HIGHLIGHT: We propose a biologically-inspired context-aware object recognition model consisting of a two-stream architecture.

1290, TITLE: SLV: Spatial Likelihood Voting for Weakly Supervised Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_SLV_Spatial_Likelihood_Voting_for_Weakly_Supervised_Object_Detection_CVPR_2020_paper.html
AUTHORS: Ze Chen, Zhihang Fu, Rongxin Jiang, Yaowu Chen, Xian-Sheng Hua
HIGHLIGHT: In this paper, we propose a spatial likelihood voting (SLV) module to converge the proposal localizing process without any bounding box annotations.

1291, TITLE: Universal Weighting Metric Learning for Cross-Modal Matching
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Universal_Weighting_Metric_Learning_for_Cross-Modal_Matching_CVPR_2020_paper.html
AUTHORS: Jiwei Wei, Xing Xu, Yang Yang, Yanli Ji, Zheng Wang, Heng Tao Shen
HIGHLIGHT: To address this problem, we propose a simple and interpretable universal weighting framework for cross-modal matching, which provides a tool to analyze the interpretability of various loss functions.

1292, TITLE: IDA-3D: Instance-Depth-Aware 3D Object Detection From Stereo Vision for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Peng_IDA-3D_Instance-Depth-Aware_3D_Object_Detection_From_Stereo_Vision_for_Autonomous_CVPR_2020_paper.html
AUTHORS: Wanli Peng, Hao Pan, He Liu, Yi Sun
HIGHLIGHT: Considering more general scenes, where there is no LiDAR data in the 3D datasets, we propose a 3D object detection approach from stereo vision which does not rely on LiDAR data either as input or as supervision in training, but solely takes RGB images with corresponding annotated 3D bounding boxes as training data.

1293, TITLE: Label Decoupling Framework for Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Label_Decoupling_Framework_for_Salient_Object_Detection_CVPR_2020_paper.html
AUTHORS: Jun Wei, Shuhui Wang, Zhe Wu, Chi Su, Qingming Huang, Qi Tian
HIGHLIGHT: To address this problem, we propose a label decoupling framework (LDF) which consists of a label decoupling (LD) procedure and a feature interaction network (FIN).

1294, TITLE: Transform and Tell: Entity-Aware News Image Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Tran_Transform_and_Tell_Entity-Aware_News_Image_Captioning_CVPR_2020_paper.html
AUTHORS: Alasdair Tran, Alexander Mathews, Lexing Xie
HIGHLIGHT: We propose an end-to-end model which generates captions for images embedded in news articles.

1295, TITLE: HAMBox: Delving Into Mining High-Quality Anchors on Face Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_HAMBox_Delving_Into_Mining_High-Quality_Anchors_on_Face_Detection_CVPR_2020_paper.html
AUTHORS: Yang Liu, Xu Tang, Junyu Han, Jingtuo Liu, Dinger Rui, Xiang Wu
HIGHLIGHT: In this paper, we propose an Online High-quality Anchor Mining Strategy (HAMBox), which explicitly helps outer faces compensate with high-quality anchors.

1296, TITLE: Hierarchical Feature Embedding for Attribute Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Hierarchical_Feature_Embedding_for_Attribute_Recognition_CVPR_2020_paper.html
AUTHORS: Jie Yang, Jiarou Fan, Yiru Wang, Yige Wang, Weihao Gan, Lin Liu, Wei Wu
HIGHLIGHT: To address this problem, we propose a hierarchical feature embedding (HFE) framework, which learns a fine-grained feature embedding by combining attribute and ID information.

1297, TITLE: Squeeze-and-Attention Networks for Semantic Segmentation

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhong_Squeeze-and-Attention_Networks_for_Semantic_Segmentation_CVPR_2020_paper.html
AUTHORS: Zilong Zhong, Zhong Qiu Lin, Rene Bidart, Xiaodan Hu, Ibrahim Ben Daya, Zhifeng Li, Wei-Shi Zheng, Jonathan Li, Alexander Wong
HIGHLIGHT: In this paper, we propose a novel squeeze-and-attention network (SANet) architecture that leverages an effective squeeze-and-attention (SA) module to account for two distinctive characteristics of segmentation: i) pixel-group attention, and ii) pixel-wise prediction.

1298, TITLE: Context R-CNN: Long Term Temporal Context for Per-Camera Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Beery_Context_R-CNN_Long_Term_Temporal_Context_for_Per-Camera_Object_Detection_CVPR_2020_paper.html
AUTHORS: Sara Beery, Guanhang Wu, Vivek Rathod, Ronny Votel, Jonathan Huang
HIGHLIGHT: In this paper we propose a method that leverages temporal context from the unlabeled frames of a novel camera to improve performance at that camera.

1299, TITLE: Mixture Dense Regression for Object Detection and Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Varamesh_Mixture_Dense_Regression_for_Object_Detection_and_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Ali Varamesh, Tinne Tuytelaars
HIGHLIGHT: To this end, we devise a framework for spatial regression using mixture density networks.

1300, TITLE: Syntax-Aware Action Targeting for Video Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Syntax-Aware_Action_Targeting_for_Video_Captioning_CVPR_2020_paper.html
AUTHORS: Qi Zheng, Chaoyue Wang, Dacheng Tao
HIGHLIGHT: Specifically, we propose a Syntax-Aware Action Targeting (SAAT) module that firstly builds a self-attended scene representation to draw global dependence among multiple objects within a scene, and then decodes the visually-related syntax components by setting different queries.

1301, TITLE: Learning Visual Emotion Representations From Web Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Learning_Visual_Emotion_Representations_From_Web_Data_CVPR_2020_paper.html
AUTHORS: Zijun Wei, Jianming Zhang, Zhe Lin, Joon-Young Lee, Niranjan Balasubramanian, Minh Hoai, Dimitris Samaras
HIGHLIGHT: We present a scalable approach for learning powerful visual features for emotion recognition.

1302, TITLE: The Edge of Depth: Explicit Constraints Between Segmentation and Depth
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_The_Edge_of_Depth_Explicit_Constraints_Between_Segmentation_and_Depth_CVPR_2020_paper.html
AUTHORS: Shengjie Zhu, Garrick Brazil, Xiaoming Liu
HIGHLIGHT: In this work we study the mutual benefits of two common computer vision tasks, self-supervised depth estimation and semantic segmentation from images.

1303, TITLE: A Context-Aware Loss Function for Action Spotting in Soccer Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Cioppa_A_Context-Aware_Loss_Function_for_Action_Spotting_in_Soccer_Videos_CVPR_2020_paper.html
AUTHORS: Anthony Cioppa, Adrien Deliege, Silvio Giancola, Bernard Ghanem, Marc Van Droogenbroeck, Rikke Gade, Thomas B. Moeslund
HIGHLIGHT: In this paper, we propose a novel loss function that specifically considers the temporal context naturally present around each action, rather than focusing on the single annotated frame to spot.

1304, TITLE: Towards Learning a Generic Agent for Vision-and-Language Navigation via Pre-Training
http://openaccess.thecvf.com/content_CVPR_2020/html/Hao_Towards_Learning_a_Generic_Agent_for_Vision-and-Language_Navigation_via_Pre-Training_CVPR_2020_paper.html
AUTHORS: Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, Jianfeng Gao
HIGHLIGHT: In this paper, we present the first pre-training and fine-tuning paradigm for vision-and-language navigation (VLN) tasks.

1305, TITLE: Video Instance Segmentation Tracking With a Modified VAE Architecture
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Video_Instance_Segmentation_Tracking_With_a_Modified_VAE_Architecture_CVPR_2020_paper.html
AUTHORS: Chung-Ching Lin, Ying Hung, Rogerio Feris, Linglin He

HIGHLIGHT: We propose a modified variational autoencoder (VAE) architecture built on top of Mask R-CNN for instance-level video segmentation and tracking.

1306, TITLE: Deformation-Aware Unpaired Image Translation for Pose Estimation on Laboratory Animals
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Deformation-Aware_Unpaired_Image_Translation_for_Pose_Estimation_on_Laboratory_Animals_CVPR_2020_paper.html
AUTHORS: Siyuan Li, Semih Gunel, Mirela Ostrek, Pavan Ramdya, Pascal Fua, Helge Rhodin
HIGHLIGHT: Our goal is to capture the pose of real animals using synthetic training examples, without using any manual supervision.

1307, TITLE: ZeroQ: A Novel Zero Shot Quantization Framework
http://openaccess.thecvf.com/content_CVPR_2020/html/Cai_ZeroQ_A_Novel_Zero_Shot_Quantization_Framework_CVPR_2020_paper.html
AUTHORS: Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W. Mahoney, Kurt Keutzer
HIGHLIGHT: Here, we propose \OURS, a novel zero-shot quantization framework to address this.

1308, TITLE: Disparity-Aware Domain Adaptation in Stereo Image Restoration
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Disparity-Aware_Domain_Adaptation_in_Stereo_Image_Restoration_CVPR_2020_paper.html
AUTHORS: Bo Yan, Chenxi Ma, Bahetiyaer Bare, Weimin Tan, Steven C. H. Hoi
HIGHLIGHT: Towards this end, this paper analyses how to effectively explore disparity information, and proposes a unified stereo image restoration framework.

1309, TITLE: Offset Bin Classification Network for Accurate Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Qiu_Offset_Bin_Classification_Network_for_Accurate_Object_Detection_CVPR_2020_paper.html
AUTHORS: Heqian Qiu, Hongliang Li, Qingbo Wu, Hengcan Shi
HIGHLIGHT: In this paper, we propose an offset bin classification network optimized with cross entropy loss to predict more accurate offsets.

1310, TITLE: TBT: Targeted Neural Network Attack With Bit Trojan
http://openaccess.thecvf.com/content_CVPR_2020/html/Rakin_TBT_Targeted_Neural_Network_Attack_With_Bit_Trojan_CVPR_2020_paper.html
AUTHORS: Adnan Siraj Rakin, Zhezhi He, Deliang Fan
HIGHLIGHT: In this work, for the first time, we propose a novel Targeted Bit Trojan(TBT) method, which can insert a targeted neural Trojan into a DNN through bit-flip attack.

1311, TITLE: Maintaining Discrimination and Fairness in Class Incremental Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Maintaining_Discrimination_and_Fairness_in_Class_Incremental_Learning_CVPR_2020_paper.html
AUTHORS: Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, Shu-Tao Xia
HIGHLIGHT: In this paper, we propose a simple and effective solution motivated by the aforementioned observations to address catastrophic forgetting.

1312, TITLE: Background Data Resampling for Outlier-Aware Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Background_Data_Resampling_for_Outlier-Aware_Classification_CVPR_2020_paper.html
AUTHORS: Yi Li, Nuno Vasconcelos
HIGHLIGHT: The problem of learning an image classifier that allows detection of out-of-distribution (OOD) examples, with the help of auxiliary background datasets, is studied.

1313, TITLE: STEFANN: Scene Text Editor Using Font Adaptive Neural Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Roy_STEFANN_Scene_Text_Editor_Using_Font_Adaptive_Neural_Network_CVPR_2020_paper.html
AUTHORS: Prasun Roy, Saumik Bhattacharya, Subhankar Ghosh, Umapada Pal
HIGHLIGHT: In this paper, we propose a method to modify text in an image at character-level.

1314, TITLE: Geometry and Learning Co-Supported Normal Estimation for Unstructured Point Cloud
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_Geometry_and_Learning_Co-Supported_Normal_Estimation_for_Unstructured_Point_Cloud_CVPR_2020_paper.html

AUTHORS: Haoran Zhou, Honghua Chen, Yidan Feng, Qiong Wang, Jing Qin, Haoran Xie, Fu Lee Wang, Mingqiang Wei, Jun Wang
HIGHLIGHT: In this paper, we propose a normal estimation method for unstructured point cloud.

1315, TITLE: Sequential Motif Profiles and Topological Plots for Offline Signature Verification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zois_Sequential_Motif_Profiles_and_Topological_Plots_for_Offline_Signature_Verification_CVPR_2020_paper.html
AUTHORS: Elias N. Zois, Evangelos Zervas, Dimitrios Tsourounis, George Economou
HIGHLIGHT: In this paper, inspired by the recent use of image visibility graphs for mapping images into networks, we introduce for the first time in offline SV literature their use as a parameter free, agnostic representation for exploring global as well as local information.

1316, TITLE: Optical Flow in Dense Foggy Scenes Using Semi-Supervised Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Optical_Flow_in_Dense_Foggy_Scenes_Using_Semi-Supervised_Learning_CVPR_2020_paper.html
AUTHORS: Wending Yan, Aashish Sharma, Robby T. Tan
HIGHLIGHT: To address the problem, we introduce a semi-supervised deep learning technique that employs real fog images without optical flow ground-truths in the training process.

1317, TITLE: A Spatial RNN Codec for End-to-End Image Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_A_Spatial_RNN_Codec_for_End-to-End_Image_Compression_CVPR_2020_paper.html
AUTHORS: Chaoyi Lin, Jiabao Yao, Fangdong Chen, Li Wang
HIGHLIGHT: In this paper, we propose a fast yet effective method for end-to-end image compression by incorporating a novel spatial recurrent neural network.

1318, TITLE: Object Relational Graph With Teacher-Recommended Learning for Video Captioning
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Object_Relational_Graph_With_Teacher-Recommended_Learning_for_Video_Captioning_CVPR_2020_paper.html
AUTHORS: Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, Zheng-Jun Zha
HIGHLIGHT: In this paper, we propose a complete video captioning system including both a novel model and an effective training strategy.

1319, TITLE: MMTM: Multimodal Transfer Module for CNN Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Joze_MMTM_Multimodal_Transfer_Module_for_CNN_Fusion_CVPR_2020_paper.html
AUTHORS: Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L. Iuzzolino, Kazuhito Koishida
HIGHLIGHT: In this paper, we present a simple neural network module for leveraging the knowledge from multiple modalities in convolutional neural networks.

1320, TITLE: Generalized Zero-Shot Learning via Over-Complete Distribution
http://openaccess.thecvf.com/content_CVPR_2020/html/Keshari_Generalized_Zero-Shot_Learning_via_Over-Complete_Distribution_CVPR_2020_paper.html
AUTHORS: Rohit Keshari, Richa Singh, Mayank Vatsa
HIGHLIGHT: To learn a discriminative classifier which yields good performance in Zero-Shot Learning (ZSL) settings, we propose to generate an Over-Complete Distribution (OCD) using Conditional Variational Autoencoder (CVAE) of both seen and unseen classes.

1321, TITLE: Gait Recognition via Semi-supervised Disentangled Representation Learning to Identity and Covariate Features
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Gait_Recognition_via_Semi-supervised_Disentangled_Representation_Learning_to_Identity_and_CVPR_2020_paper.html
AUTHORS: Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, Mingwu Ren
HIGHLIGHT: We therefore propose a method of gait recognition via disentangled representation learning that considers both identity and covariate features.

1322, TITLE: Unifying Training and Inference for Panoptic Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Unifying_Training_and_Inference_for_Panoptic_Segmentation_CVPR_2020_paper.html
AUTHORS: Qizhu Li, Xiaojuan Qi, Philip H.S. Torr
HIGHLIGHT: We present an end-to-end network to bridge the gap between training and inference pipeline for panoptic segmentation, a task that seeks to partition an image into semantic regions for "stuff" and object instances for "things".

1323, TITLE: Associate-3Ddet: Perceptual-to-Conceptual Association for 3D Point Cloud Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Du_Associate-3Ddet_Perceptual-to-Conceptual_Association_for_3D_Point_Cloud_Object_Detection_CVPR_2020_paper.html
AUTHORS: Liang Du, Xiaoqing Ye, Xiao Tan, Jianfeng Feng, Zhenbo Xu, Errui Ding, Shilei Wen
HIGHLIGHT: In this paper, we innovatively propose a domain adaptation like approach to enhance the robustness of the feature representation.

1324, TITLE: Interactive Image Segmentation With First Click Attention
http://openaccess.thecvf.com/content_CVPR_2020/html/Lin_Interactive_Image_Segmentation_With_First_Click_Attention_CVPR_2020_paper.html
AUTHORS: Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, Shao-Ping Lu
HIGHLIGHT: In this paper, we demonstrate the critical role of the first click about providing the location and main body information of the target object.

1325, TITLE: NETNet: Neighbor Erasing and Transferring Network for Better Single Shot Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_NETNet_Neighbor_Erasing_and_Transferring_Network_for_Better_Single_Shot_CVPR_2020_paper.html
AUTHORS: Yazhao Li, Yanwei Pang, Jianbing Shen, Jiale Cao, Ling Shao
HIGHLIGHT: With this observation, we propose a new Neighbor Erasing and Transferring (NET) mechanism to reconfigure the pyramid features and explore scale-aware features.

1326, TITLE: Scale-Equalizing Pyramid Convolution for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Scale-Equalizing_Pyramid_Convolution_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Xinjiang Wang, Shilong Zhang, Zhuoran Yu, Litong Feng, Wayne Zhang
HIGHLIGHT: Inspired by this, a convolution across the pyramid level is proposed in this study, which is termed pyramid convolution and is a modified 3-D convolution.

1327, TITLE: Learning to Cluster Faces via Confidence and Connectivity Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Learning_to_Cluster_Faces_via_Confidence_and_Connectivity_Estimation_CVPR_2020_paper.html
AUTHORS: Lei Yang, Dapeng Chen, Xiaohang Zhan, Rui Zhao, Chen Change Loy, Dahua Lin
HIGHLIGHT: In this paper, we propose a fully learnable clustering framework without requiring a large number of overlapped subgraphs.

1328, TITLE: Cross-Modality Person Re-Identification With Shared-Specific Feature Transfer
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_Cross-Modality_Person_Re-Identification_With_Shared-Specific_Feature_Transfer_CVPR_2020_paper.html
AUTHORS: Yan Lu, Yue Wu, Bin Liu, Tianzhu Zhang, Baopu Li, Qi Chu, Nenghai Yu
HIGHLIGHT: In this paper, we tackle the above limitation by proposing a novel cross-modality shared-specific feature transfer algorithm (termed cm-SSFT) to explore the potential of both the modality-shared information and the modality-specific characteristics to boost the reidentification performance.

1329, TITLE: DPGN: Distribution Propagation Graph Network for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_DPGN_Distribution_Propagation_Graph_Network_for_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Ling Yang, Liangliang Li, Zilun Zhang, Xinyu Zhou, Erjin Zhou, Yu Liu
HIGHLIGHT: We propose a novel approach named distribution propagation graph network (DPGN) for few-shot learning.

1330, TITLE: Density-Aware Graph for Deep Semi-Supervised Visual Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Density-Aware_Graph_for_Deep_Semi-Supervised_Visual_Recognition_CVPR_2020_paper.html
AUTHORS: Suichan Li, Bin Liu, Dongdong Chen, Qi Chu, Lu Yuan, Nenghai Yu
HIGHLIGHT: Motivated by these limitations, this paper proposes to solve the SSL problem by building a novel density-aware graph, based on which the neighborhood information can be easily leveraged and the feature learning and label propagation can also be trained in an end-to-end way.

1331, TITLE: Unsupervised Multi-Modal Image Registration via Geometry Preserving Image-to-Image Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Arar_Unsupervised_Multi-Modal_Image_Registration_via_Geometry_Preserving_Image-to-Image_Translation_CVPR_2020_paper.html

AUTHORS: Moab Arar, Yiftach Ginger, Dov Danon, Amit H. Bermano, Daniel Cohen-Or
HIGHLIGHT: In this work, we bypass the difficulties of developing cross-modality similarity measures, by training an image-to-image translation network on the two input modalities.

1332, TITLE: Binarizing MobileNet via Evolution-Based Searching
http://openaccess.thecvf.com/content_CVPR_2020/html/Phan_Binarizing_MobileNet_via_Evolution-Based_Searching_CVPR_2020_paper.html
AUTHORS: Hai Phan, Zechun Liu, Dang Huynh, Marios Savvides, Kwang-Ting Cheng, Zhiqiang Shen
HIGHLIGHT: In this paper, we propose a use of evolutionary search to facilitate the construction and training scheme when binarizing MobileNet, a compact network with separable depth-wise convolution.

1333, TITLE: Temporal-Context Enhanced Detection of Heavily Occluded Pedestrians
http://openaccess.thecvf.com/content_CVPR_2020/html/Wu_Temporal-Context_Enhanced_Detection_of_Heavily_Occluded_Pedestrians_CVPR_2020_paper.html
AUTHORS: Jialian Wu, Chunluan Zhou, Ming Yang, Qian Zhang, Yuan Li, Junsong Yuan
HIGHLIGHT: In this paper, we exploit the local temporal context of pedestrians in videos and propose a tube feature aggregation network (TFAN) aiming at enhancing pedestrian detectors against severe occlusions.

1334, TITLE: Orderless Recurrent Models for Multi-Label Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Yazici_Orderless_Recurrent_Models_for_Multi-Label_Classification_CVPR_2020_paper.html
AUTHORS: Vacit Oguz Yazici, Abel Gonzalez-Garcia, Arnau Ramisa, Bartlomiej Twardowski, Joost van de Weijer
HIGHLIGHT: Therefore, in this paper, we propose ways to dynamically order the ground truth labels with the predicted label sequence.

1335, TITLE: Gold Seeker: Information Gain From Policy Distributions for Goal-Oriented Vision-and-Langauge Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Abbasnejad_Gold_Seeker_Information_Gain_From_Policy_Distributions_for_Goal-Oriented_Vision-and-Langauge_CVPR_2020_paper.html
AUTHORS: Ehsan Abbasnejad, Iman Abbasnejad, Qi Wu, Javen Shi, Anton van den Hengel
HIGHLIGHT: We propose a reinforcement-learning approach that maintains a distribution over its internal information, thus explicitly representing the ambiguity in what it knows, and needs to know, towards achieving its goal.

1336, TITLE: Rethinking the Route Towards Weakly Supervised Object Localization
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Rethinking_the_Route_Towards_Weakly_Supervised_Object_Localization_CVPR_2020_paper.html
AUTHORS: Chen-Lin Zhang, Yun-Hao Cao, Jianxin Wu
HIGHLIGHT: In this paper, we demonstrate that weakly supervised object localization should be divided into two parts: class-agnostic object localization and object classification.

1337, TITLE: Adversarial Feature Hallucination Networks for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Adversarial_Feature_Hallucination_Networks_for_Few-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Kai Li, Yulun Zhang, Kunpeng Li, Yun Fu
HIGHLIGHT: In this paper, we propose Adversarial Feature Hallucination Networks (AFHN) which is based on conditional Wasserstein Generative Adversarial networks (cWGAN) and hallucinates diverse and discriminative features conditioned on the few labeled samples.

1338, TITLE: Conditional Gaussian Distribution Learning for Open Set Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Sun_Conditional_Gaussian_Distribution_Learning_for_Open_Set_Recognition_CVPR_2020_paper.html
AUTHORS: Xin Sun, Zhenning Yang, Chi Zhang, Keck-Voon Ling, Guohao Peng
HIGHLIGHT: In this paper, we propose a novel method, Conditional Gaussian Distribution Learning (CGDL), for open set recognition.

1339, TITLE: Connect-and-Slice: An Hybrid Approach for Reconstructing 3D Objects
http://openaccess.thecvf.com/content_CVPR_2020/html/Fang_Connect-and-Slice_An_Hybrid_Approach_for_Reconstructing_3D_Objects_CVPR_2020_paper.html
AUTHORS: Hao Fang, Florent Lafarge
HIGHLIGHT: In this paper, we address this issue with an hybrid method that successively connects and slices planes detected from 3D data.

1340, TITLE:	Attentive Weights Generation for Few Shot Learning via Information Maximization
http://openaccess.thecvf.com/content_CVPR_2020/html/Guo_Attentive_Weights_Generation_for_Few_Shot_Learning_via_Information_Maximization_CVPR_2020_paper.html
AUTHORS:	Yiluan Guo, Ngai-Man Cheung
HIGHLIGHT:	In this work, we present Attentive Weights Generation for few shot learning via Information Maximization (AWGIM), which introduces two novel contributions: i) Mutual information maximization between generated weights and data within the task; this enables the generated weights to retain information of the task and the specific query sample.

1341, TITLE:	Assessing Eye Aesthetics for Automatic Multi-Reference Eye In-Painting
http://openaccess.thecvf.com/content_CVPR_2020/html/Yan_Assessing_Eye_Aesthetics_for_Automatic_Multi-Reference_Eye_In-Painting_CVPR_2020_paper.html
AUTHORS:	Bo Yan, Qing Lin, Weimin Tan, Shili Zhou
HIGHLIGHT:	In this paper, aesthetic assessment is introduced into eye in-painting task for the first time.
We construct an eye aesthetic dataset, and train the eye aesthetic assessment network on this basis.

1342, TITLE:	PuppeteerGAN: Arbitrary Portrait Animation With Semantic-Aware Appearance Transformation
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_PuppeteerGAN_Arbitrary_Portrait_Animation_With_Semantic-Aware_Appearance_Transformation_CVPR_2020_paper.html
AUTHORS:	Zhuo Chen, Chaoyue Wang, Bo Yuan, Dacheng Tao
HIGHLIGHT:	In this paper, we devised a novel two-stage framework called PuppeteerGAN for solving these challenges.

1343, TITLE:	SEED: Semantics Enhanced Encoder-Decoder Framework for Scene Text Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Qiao_SEED_Semantics_Enhanced_Encoder-Decoder_Framework_for_Scene_Text_Recognition_CVPR_2020_paper.html
AUTHORS:	Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, Weiping Wang
HIGHLIGHT:	In this work, we propose a semantics enhanced encoder-decoder framework to robustly recognize low-quality scene texts.

1344, TITLE:	Texture and Shape Biased Two-Stream Networks for Clothing Classification and Attribute Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Texture_and_Shape_Biased_Two-Stream_Networks_for_Clothing_Classification_and_CVPR_2020_paper.html
AUTHORS:	Yuwei Zhang, Peng Zhang, Chun Yuan, Zhi Wang
HIGHLIGHT:	To this end, we propose to use two streams to enhance the extraction of shape and texture, respectively.

1345, TITLE:	Distortion Agnostic Deep Watermarking
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Distortion_Agnostic_Deep_Watermarking_CVPR_2020_paper.html
AUTHORS:	Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, Peyman Milanfar
HIGHLIGHT:	In this paper, we propose a new framework for distortion-agnostic watermarking, where the image distortion is not explicitly modeled during training.

1346, TITLE:	RMP-SNN: Residual Membrane Potential Neuron for Enabling Deeper High-Accuracy and Low-Latency Spiking Neural Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_RMP-SNN_Residual_Membrane_Potential_Neuron_for_Enabling_Deeper_High-Accuracy_and_CVPR_2020_paper.html
AUTHORS:	Bing Han, Gopalakrishnan Srinivasan, Kaushik Roy
HIGHLIGHT:	We propose ANN-SNN conversion using "soft reset" spiking neuron model, referred to as Residual Membrane Potential (RMP) spiking neuron, which retains the "residual" membrane potential above threshold at the firing instants.

1347, TITLE:	BFBox: Searching Face-Appropriate Backbone and Feature Pyramid Network for Face Detector
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_BFBox_Searching_Face-Appropriate_Backbone_and_Feature_Pyramid_Network_for_Face_CVPR_2020_paper.html
AUTHORS:	Yang Liu, Xu Tang
HIGHLIGHT:	To resolve this, the success of Neural Archi-tecture Search (NAS) inspires us to search face-appropriate backbone and featrue pyramid network (FPN) architecture.

1348, TITLE:	PFCNN: Convolutional Neural Networks on 3D Surfaces Using Parallel Frames
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_PFCNN_Convolutional_Neural_Networks_on_3D_Surfaces_Using_Parallel_Frames_CVPR_2020_paper.html
AUTHORS:	Yuqi Yang, Shilin Liu, Hao Pan, Yang Liu, Xin Tong
HIGHLIGHT:	We use parallel frames on surface to define PFCNNs that enable effective feature learning on surface meshes by mimicking standard convolutions faithfully.

1349, TITLE:		iTAML: An Incremental Task-Agnostic Meta-learning Approach
http://openaccess.thecvf.com/content_CVPR_2020/html/Rajasegaran_iTAML_An_Incremental_Task-Agnostic_Meta-learning_Approach_CVPR_2020_paper.html
AUTHORS:		Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Mubarak Shah
HIGHLIGHT:		In this pursuit, we introduce a novel meta-learning approach that seeks to maintain an equilibrium between all the encountered tasks.

1350, TITLE:		Optimal least-squares solution to the hand-eye calibration problem
http://openaccess.thecvf.com/content_CVPR_2020/html/Dekel_Optimal_least-squares_solution_to_the_hand-eye_calibration_problem_CVPR_2020_paper.html
AUTHORS:		Amit Dekel, Linus Harenstam-Nielsen, Sergio Caccamo
HIGHLIGHT:		We propose a least-squares formulation to the noisy hand-eye calibration problem using dual-quaternions, and introduce efficient algorithms to find the exact optimal solution, based on analytic properties of the problem, avoiding non-linear optimization.

1351, TITLE:		MnasFPN: Learning Latency-Aware Pyramid Architecture for Object Detection on Mobile Devices
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_MnasFPN_Learning_Latency-Aware_Pyramid_Architecture_for_Object_Detection_on_Mobile_CVPR_2020_paper.html
AUTHORS:		Bo Chen, Golnaz Ghiasi, Hanxiao Liu, Tsung-Yi Lin, Dmitry Kalenichenko, Hartwig Adam, Quoc V. Le
HIGHLIGHT:		We propose MnasFPN, a mobile-friendly search space for the detection head, and combine it with latency-aware architecture search to produce efficient object detection models.

1352, TITLE:		VSGNet: Spatial Attention Network for Detecting Human Object Interactions Using Graph Convolutions
http://openaccess.thecvf.com/content_CVPR_2020/html/Ulutan_VSGNet_Spatial_Attention_Network_for_Detecting_Human_Object_Interactions_Using_CVPR_2020_paper.html
AUTHORS:		Oytun Ulutan, A S M Iftekhar, B. S. Manjunath
HIGHLIGHT:		VSGNet extracts visual features from the human-object pairs, refines the features with spatial configurations of the pair, and utilizes the structural connections between the pair via graph convolutions.

1353, TITLE:		End-to-End Camera Calibration for Broadcast Videos
http://openaccess.thecvf.com/content_CVPR_2020/html/Sha_End-to-End_Camera_Calibration_for_Broadcast_Videos_CVPR_2020_paper.html
AUTHORS:		Long Sha, Jennifer Hobbs, Panna Felsen, Xinyu Wei, Patrick Lucey, Sujoy Ganguly
HIGHLIGHT:		In this paper, we propose an end-to-end approach for single moving camera calibration across challenging scenarios in sports.

1354, TITLE:		Regularizing CNN Transfer Learning With Randomised Regression
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhong_Regularizing_CNN_Transfer_Learning_With_Randomised_Regression_CVPR_2020_paper.html
AUTHORS:		Yang Zhong, Atsuto Maki
HIGHLIGHT:		This paper is about regularizing deep convolutional networks (CNNs) based on an adaptive framework for transfer learning with limited training data in the target domain.

1355, TITLE:		KeypointNet: A Large-Scale 3D Keypoint Dataset Aggregated From Numerous Human Annotations
http://openaccess.thecvf.com/content_CVPR_2020/html/You_KeypointNet_A_Large-Scale_3D_Keypoint_Dataset_Aggregated_From_Numerous_Human_CVPR_2020_paper.html
AUTHORS:		Yang You, Yujing Lou, Chengkun Li, Zhoujun Cheng, Liangwei Li, Lizhuang Ma, Cewu Lu, Weiming Wang
HIGHLIGHT:		To handle the inconsistency between annotations from different people, we propose a novel method to aggregate these keypoints automatically, through minimization of a fidelity loss.

1356, TITLE:		Hierarchical Clustering With Hard-Batch Triplet Loss for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2020/html/Zeng_Hierarchical_Clustering_With_Hard-Batch_Triplet_Loss_for_Person_Re-Identification_CVPR_2020_paper.html
AUTHORS:		Kaiwei Zeng, Munan Ning, Yaohua Wang, Yang Guo
HIGHLIGHT:		In order to improve the quality of pseudo labels in existing methods, we propose the HCT method which combines hierarchical clustering with hard-batch triplet loss.

1357, TITLE:		Joint Semantic Segmentation and Boundary Detection Using Iterative Pyramid Contexts
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhen_Joint_Semantic_Segmentation_and_Boundary_Detection_Using_Iterative_Pyramid_Contexts_CVPR_2020_paper.html

AUTHORS: Mingmin Zhen, Jinglu Wang, Lei Zhou, Shiwei Li, Tianwei Shen, Jiaxiang Shang, Tian Fang, Long Quan
HIGHLIGHT: In this paper, we present a joint multi-task learning framework for semantic segmentation and boundary detection.

1358, TITLE: Attention-Guided Hierarchical Structure Aggregation for Image Matting
http://openaccess.thecvf.com/content_CVPR_2020/html/Qiao_Attention-Guided_Hierarchical_Structure_Aggregation_for_Image_Matting_CVPR_2020_paper.html
AUTHORS: Yu Qiao, Yuhao Liu, Xin Yang, Dongsheng Zhou, Mingliang Xu, Qiang Zhang, Xiaopeng Wei
HIGHLIGHT: In this paper, we propose an end-to-end Hierarchical Attention Matting Network (HAttMatting), which can predict the better structure of alpha mattes from single RGB images without additional input.

1359, TITLE: MetaFuse: A Pre-trained Fusion Model for Human Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Xie_MetaFuse_A_Pre-trained_Fusion_Model_for_Human_Pose_Estimation_CVPR_2020_paper.html
AUTHORS: Rongchang Xie, Chunyu Wang, Yizhou Wang
HIGHLIGHT: In this work, we introduce MetaFuse, a pre-trained fusion model learned from a large number of cameras in the Panoptic dataset.

1360, TITLE: Prior Guided GAN Based Semantic Inpainting
http://openaccess.thecvf.com/content_CVPR_2020/html/Lahiri_Prior_Guided_GAN_Based_Semantic_Inpainting_CVPR_2020_paper.html
AUTHORS: Avisek Lahiri, Arnav Kumar Jain, Sanskar Agrawal, Pabitra Mitra, Prabir Kumar Biswas
HIGHLIGHT: In this paper, going against the general trend, we focus on the second paradigm of inpainting and address both of its mentioned problems.

1361, TITLE: Weakly Supervised Semantic Point Cloud Segmentation: Towards 10x Fewer Labels
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Weakly_Supervised_Semantic_Point_Cloud_Segmentation_Towards_10x_Fewer_Labels_CVPR_2020_paper.html
AUTHORS: Xun Xu, Gim Hee Lee
HIGHLIGHT: In this work, we propose a weakly supervised point cloud segmentation approach which requires only a tiny fraction of points to be labelled in the training stage.

1362, TITLE: Physically Realizable Adversarial Examples for LiDAR Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Tu_Physically_Realizable_Adversarial_Examples_for_LiDAR_Object_Detection_CVPR_2020_paper.html
AUTHORS: James Tu, Mengye Ren, Sivabalan Manivasagam, Ming Liang, Bin Yang, Richard Du, Frank Cheng, Raquel Urtasun
HIGHLIGHT: In this paper, we address this issue and present a method to generate universal 3D adversarial objects to fool LiDAR detectors.

1363, TITLE: Combating Noisy Labels by Agreement: A Joint Training Method with Co-Regularization
http://openaccess.thecvf.com/content_CVPR_2020/html/Wei_Combating_Noisy_Labels_by_Agreement_A_Joint_Training_Method_with_CVPR_2020_paper.html
AUTHORS: Hongxin Wei, Lei Feng, Xiangyu Chen, Bo An
HIGHLIGHT: In this paper, we start from a different perspective and propose a robust learning paradigm called JoCoR, which aims to reduce the diversity of two networks during training.

1364, TITLE: Light-weight Calibrator: A Separable Component for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Ye_Light-weight_Calibrator_A_Separable_Component_for_Unsupervised_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Shaokai Ye, Kailu Wu, Mu Zhou, Yunfei Yang, Sia Huat Tan, Kaidi Xu, Jiebo Song, Chenglong Bao, Kaisheng Ma
HIGHLIGHT: In this work, instead of training a classifier to adapt to the target domain, we use a separable component called data calibrator to help the fixed source classifier recover discrimination power in the target domain, while preserving the source domain's performance.

1365, TITLE: Learn to Augment: Joint Data Augmentation and Network Optimization for Text Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Learn_to_Augment_Joint_Data_Augmentation_and_Network_Optimization_for_CVPR_2020_paper.html
AUTHORS: Canjie Luo, Yuanzhi Zhu, Lianwen Jin, Yongpan Wang
HIGHLIGHT: In this paper, we propose a new method for text image augmentation.

1366, TITLE: Learning Selective Self-Mutual Attention for RGB-D Saliency Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Learning_Selective_Self-Mutual_Attention_for_RGB-D_Saliency_Detection_CVPR_2020_paper.html
AUTHORS: Nian Liu, Ni Zhang, Junwei Han
HIGHLIGHT: In this paper, we propose to fuse attention learned in both modalities.

1367, TITLE: Cross-domain Object Detection through Coarse-to-Fine Feature Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Zheng_Cross-domain_Object_Detection_through_Coarse-to-Fine_Feature_Adaptation_CVPR_2020_paper.html
AUTHORS: Yangtao Zheng, Di Huang, Songtao Liu, Yunhong Wang
HIGHLIGHT: To address such an issue, this paper proposes a novel coarse-to-fine feature adaptation approach to cross-domain object detection.

1368, TITLE: Estimating Low-Rank Region Likelihood Maps
http://openaccess.thecvf.com/content_CVPR_2020/html/Csurka_Estimating_Low-Rank_Region_Likelihood_Maps_CVPR_2020_paper.html
AUTHORS: Gabriela Csurka, Zoltan Kato, Andor Juhasz, Martin Humenberger
HIGHLIGHT: Herein, we propose a novel self-supervised low-rank region detection deep network that predicts a low-rank likelihood map from an image.

1369, TITLE: Neural Head Reenactment with Latent Pose Descriptors
http://openaccess.thecvf.com/content_CVPR_2020/html/Burkov_Neural_Head_Reenactment_with_Latent_Pose_Descriptors_CVPR_2020_paper.html
AUTHORS: Egor Burkov, Igor Pasechnik, Artur Grigorev, Victor Lempitsky
HIGHLIGHT: We propose a neural head reenactment system, which is driven by a latent pose representation and is capable of predicting the foreground segmentation alongside the RGB image.

1370, TITLE: Learning Individual Speaking Styles for Accurate Lip to Speech Synthesis
http://openaccess.thecvf.com/content_CVPR_2020/html/Prajwal_Learning_Individual_Speaking_Styles_for_Accurate_Lip_to_Speech_Synthesis_CVPR_2020_paper.html
AUTHORS: K R Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, C.V. Jawahar
HIGHLIGHT: In this work, we explore the task of lip to speech synthesis, i.e., learning to generate natural speech given only the lip movements of a speaker.

1371, TITLE: Self-Supervised Learning of Video-Induced Visual Invariances
http://openaccess.thecvf.com/content_CVPR_2020/html/Tschannen_Self-Supervised_Learning_of_Video-Induced_Visual_Invariances_CVPR_2020_paper.html
AUTHORS: Michael Tschannen, Josip Djolonga, Marvin Ritter, Aravindh Mahendran, Neil Houlsby, Sylvain Gelly, Mario Lucic
HIGHLIGHT: We propose a general framework for self-supervised learning of transferable visual representations based on Video-Induced Visual Invariances (VIVI).

1372, TITLE: Two-Stage Peer-Regularized Feature Recombination for Arbitrary Image Style Transfer
http://openaccess.thecvf.com/content_CVPR_2020/html/Svoboda_Two-Stage_Peer-Regularized_Feature_Recombination_for_Arbitrary_Image_Style_Transfer_CVPR_2020_paper.html
AUTHORS: Jan Svoboda, Asha Anoosheh, Christian Osendorfer, Jonathan Masci
HIGHLIGHT: This paper introduces a neural style transfer model to generate a stylized image conditioning on a set of examples describing the desired style.

1373, TITLE: MINA: Convex Mixed-Integer Programming for Non-Rigid Shape Alignment
http://openaccess.thecvf.com/content_CVPR_2020/html/Bernard_MINA_Convex_Mixed-Integer_Programming_for_Non-Rigid_Shape_Alignment_CVPR_2020_paper.html
AUTHORS: Florian Bernard, Zeeshan Khan Suri, Christian Theobalt
HIGHLIGHT: To this end, we propose a novel shape deformation model based on an efficient low-dimensional discrete model, so that finding a globally optimal solution is tractable in (most) practical cases.

1374, TITLE: Improving One-Shot NAS by Suppressing the Posterior Fading
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Improving_One-Shot_NAS_by_Suppressing_the_Posterior_Fading_CVPR_2020_paper.html
AUTHORS: Xiang Li, Chen Lin, Chuming Li, Ming Sun, Wei Wu, Junjie Yan, Wanli Ouyang

HIGHLIGHT: In this paper, we analyse existing weight sharing one-shot NAS approaches from a Bayesian point of view and identify the Posterior Fading problem, which compromises the effectiveness of shared weights.

1375, TITLE: Incremental Few-Shot Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Perez-Rua_Incremental_Few-Shot_Object_Detection_CVPR_2020_paper.html
AUTHORS: Juan-Manuel Perez-Rua, Xiatian Zhu, Timothy M. Hospedales, Tao Xiang
HIGHLIGHT: We present the first study aiming to go beyond these limitations by considering the Incremental Few-Shot Detection (iFSD) problem setting, where new classes must be registered incrementally (without revisiting base classes) and with few examples.

1376, TITLE: Synthetic Learning: Learn From Distributed Asynchronized Discriminator GAN Without Sharing Medical Image Data
http://openaccess.thecvf.com/content_CVPR_2020/html/Chang_Synthetic_Learning_Learn_From_Distributed_Asynchronized_Discriminator_GAN_Without_Sharing_CVPR_2020_paper.html
AUTHORS: Qi Chang, Hui Qu, Yikai Zhang, Mert Sabuncu, Chao Chen, Tong Zhang, Dimitris N. Metaxas
HIGHLIGHT: In this paper, we propose a data privacy-preserving and communication efficient distributed GAN learning framework named Distributed Asynchronized Discriminator GAN (AsynDGAN).

1377, TITLE: Exploring Category-Agnostic Clusters for Open-Set Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Pan_Exploring_Category-Agnostic_Clusters_for_Open-Set_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Yingwei Pan, Ting Yao, Yehao Li, Chong-Wah Ngo, Tao Mei
HIGHLIGHT: In this paper, we address this problem by augmenting the state-of-the-art domain adaptation technique, Self-Ensembling, with category-agnostic clusters in target domain.

1378, TITLE: Regularizing Class-Wise Predictions via Self-Knowledge Distillation
http://openaccess.thecvf.com/content_CVPR_2020/html/Yun_Regularizing_Class-Wise_Predictions_via_Self-Knowledge_Distillation_CVPR_2020_paper.html
AUTHORS: Sukmin Yun, Jongjin Park, Kimin Lee, Jinwoo Shin
HIGHLIGHT: To mitigate the issue, we propose a new regularization method that penalizes the predictive distribution between similar samples.

1379, TITLE: Hierarchical Graph Attention Network for Visual Relationship Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Mi_Hierarchical_Graph_Attention_Network_for_Visual_Relationship_Detection_CVPR_2020_paper.html
AUTHORS: Li Mi, Zhenzhong Chen
HIGHLIGHT: In this work, a Hierarchical Graph Attention Network (HGAT) is proposed to capture the dependencies on both object-level and triplet-level.

1380, TITLE: M2m: Imbalanced Classification via Major-to-Minor Translation
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_M2m_Imbalanced_Classification_via_Major-to-Minor_Translation_CVPR_2020_paper.html
AUTHORS: Jaehyung Kim, Jongheon Jeong, Jinwoo Shin
HIGHLIGHT: In this paper, we explore a novel yet simple way to alleviate this issue by augmenting less-frequent classes via translating samples (e.g., images) from more-frequent classes.

1381, TITLE: CenterMask: Real-Time Anchor-Free Instance Segmentation
http://openaccess.thecvf.com/content_CVPR_2020/html/Lee_CenterMask_Real-Time_Anchor-Free_Instance_Segmentation_CVPR_2020_paper.html
AUTHORS: Youngwan Lee, Jongyoul Park
HIGHLIGHT: We propose a simple yet efficient anchor-free instance segmentation, called CenterMask, that adds a novel spatial attention-guided mask (SAG-Mask) branch to anchor-free one stage object detector (FCOS) in the same vein with Mask R-CNN.

1382, TITLE: Multi-Path Learning for Object Pose Estimation Across Domains
http://openaccess.thecvf.com/content_CVPR_2020/html/Sundermeyer_Multi-Path_Learning_for_Object_Pose_Estimation_Across_Domains_CVPR_2020_paper.html
AUTHORS: Martin Sundermeyer, Maximilian Durner, En Yen Puang, Zoltan-Csaba Marton, Narunas Vaskevicius, Kai O. Arras, Rudolph Triebel
HIGHLIGHT: We introduce a scalable approach for object pose estimation trained on simulated RGB views of multiple 3D models together.

1383, TITLE: Incremental Learning in Online Scenario
http://openaccess.thecvf.com/content_CVPR_2020/html/He_Incremental_Learning_in_Online_Scenario_CVPR_2020_paper.html
AUTHORS: Jiangpeng He, Runyu Mao, Zeman Shao, Fengqing Zhu
HIGHLIGHT: In this paper, we propose an incremental learning framework that can work in the challenging online learning scenario and handle both new classes data and new observations of old classes.

1384, TITLE: Enhanced Transport Distance for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Enhanced_Transport_Distance_for_Unsupervised_Domain_Adaptation_CVPR_2020_paper.html
AUTHORS: Mengxue Li, Yi-Ming Zhai, You-Wei Luo, Peng-Fei Ge, Chuan-Xian Ren
HIGHLIGHT: In this work, we propose an enhanced transport distance (ETD) for UDA.

1385, TITLE: TESA: Tensor Element Self-Attention via Matricization
http://openaccess.thecvf.com/content_CVPR_2020/html/Babiloni_TESA_Tensor_Element_Self-Attention_via_Matricization_CVPR_2020_paper.html
AUTHORS: Francesca Babiloni, Ioannis Marras, Gregory Slabaugh, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we introduce a new method, called Tensor Element Self-Attention (TESA) that generalizes such work to capture interdependencies along all dimensions of the tensor using matricization.

1386, TITLE: Training a Steerable CNN for Guidewire Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Training_a_Steerable_CNN_for_Guidewire_Detection_CVPR_2020_paper.html
AUTHORS: Donghang Li, Adrian Barbu
HIGHLIGHT: In this paper, we present a steerable Convolutional Neural Network (CNN), which is a Fully Convolutional Neural Network (FCNN) that can detect objects rotated by an arbitrary 2D angle, without being rotation invariant.

1387, TITLE: Superpixel Segmentation With Fully Convolutional Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_Superpixel_Segmentation_With_Fully_Convolutional_Networks_CVPR_2020_paper.html
AUTHORS: Fengting Yang, Qian Sun, Hailin Jin, Zihan Zhou
HIGHLIGHT: Inspired by an initialization strategy commonly adopted by traditional superpixel algorithms, we present a novel method that employs a simple fully convolutional network to predict superpixels on a regular image grid.

1388, TITLE: SharinGAN: Combining Synthetic and Real Data for Unsupervised Geometry Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/PNVR_SharinGAN_Combining_Synthetic_and_Real_Data_for_Unsupervised_Geometry_Estimation_CVPR_2020_paper.html
AUTHORS: Koutilya PNVR, Hao Zhou, David Jacobs
HIGHLIGHT: We propose a novel method for combining synthetic and real images when training networks to determine geometric information from a single image.

1389, TITLE: Label Distribution Learning on Auxiliary Label Space Graphs for Facial Expression Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Label_Distribution_Learning_on_Auxiliary_Label_Space_Graphs_for_Facial_CVPR_2020_paper.html
AUTHORS: Shikai Chen, Jianfeng Wang, Yuedong Chen, Zhongchao Shi, Xin Geng, Yong Rui
HIGHLIGHT: To solve the problem, we propose a novel approach named Label Distribution Learning on Auxiliary Label Space Graphs(LDL-ALSG) that leverages the topological information of the labels from related but more distinct tasks, such as action unit recognition and facial landmark detection.

1390, TITLE: Deep Residual Flow for Out of Distribution Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Zisselman_Deep_Residual_Flow_for_Out_of_Distribution_Detection_CVPR_2020_paper.html
AUTHORS: Ev Zisselman, Aviv Tamar
HIGHLIGHT: In this work, we present a novel approach that improves upon the state-of-the-art by leveraging an expressive density model based on normalizing flows.

1391, TITLE: FeatureFlow: Robust Video Interpolation via Structure-to-Texture Generation
http://openaccess.thecvf.com/content_CVPR_2020/html/Gui_FeatureFlow_Robust_Video_Interpolation_via_Structure-to-Texture_Generation_CVPR_2020_paper.html
AUTHORS: Shurui Gui, Chaoyue Wang, Qihua Chen, Dacheng Tao

HIGHLIGHT: In this work, we devised a novel structure-to-texture generation framework which splits the video interpolation task into two stages: structure-guided interpolation and texture refinement.

1392, TITLE: Learning Nanoscale Motion Patterns of Vesicles in Living Cells
http://openaccess.thecvf.com/content_CVPR_2020/html/Sekh_Learning_Nanoscale_Motion_Patterns_of_Vesicles_in_Living_Cells_CVPR_2020_paper.html
AUTHORS: Arif Ahmed Sekh, Ida Sundvor Opstad, Asa Birna Birgisdottir, Truls Myrmel, Balpreet Singh Ahluwalia, Krishna Agarwal, Dilip K. Prasad
HIGHLIGHT: We propose an integrative approach, built upon physics based simulations, nanoscopy algorithms, and shallow residual attention network to make it possible for the first time to analysis sub-resolution motion patterns in vesicles that may also be of sub-resolution diameter.

1393, TITLE: Improving Action Segmentation via Graph-Based Temporal Reasoning
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Improving_Action_Segmentation_via_Graph-Based_Temporal_Reasoning_CVPR_2020_paper.html
AUTHORS: Yifei Huang, Yusuke Sugano, Yoichi Sato
HIGHLIGHT: In this paper, we propose a network module called Graph-based Temporal Reasoning Module (GTRM) that can be built on top of existing action segmentation models to learn the relation of multiple action segments in various time spans.

1394, TITLE: Episode-Based Prototype Generating Network for Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Yu_Episode-Based_Prototype_Generating_Network_for_Zero-Shot_Learning_CVPR_2020_paper.html
AUTHORS: Yunlong Yu, Zhong Ji, Jungong Han, Zhongfei Zhang
HIGHLIGHT: We introduce a simple yet effective episode-based training framework for zero-shot learning (ZSL), where the learning system requires to recognize unseen classes given only the corresponding class semantics.

1395, TITLE: Learning to Segment the Tail
http://openaccess.thecvf.com/content_CVPR_2020/html/Hu_Learning_to_Segment_the_Tail_CVPR_2020_paper.html
AUTHORS: Xinting Hu, Yi Jiang, Kaihua Tang, Jingyuan Chen, Chunyan Miao, Hanwang Zhang
HIGHLIGHT: We propose a "divide&conquer" strategy for the challenging LVIS task: divide the whole data into balanced parts and then apply incremental learning to conquer each one.

1396, TITLE: Learning to Evaluate Perception Models Using Planner-Centric Metrics
http://openaccess.thecvf.com/content_CVPR_2020/html/Philion_Learning_to_Evaluate_Perception_Models_Using_Planner-Centric_Metrics_CVPR_2020_paper.html
AUTHORS: Jonah Philion, Amlan Kar, Sanja Fidler
HIGHLIGHT: In this paper, we propose a principled metric for 3D object detection specifically for the task of self-driving.

1397, TITLE: Where, What, Whether: Multi-Modal Learning Meets Pedestrian Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Luo_Where_What_Whether_Multi-Modal_Learning_Meets_Pedestrian_Detection_CVPR_2020_paper.html
AUTHORS: Yan Luo, Chongyang Zhang, Muming Zhao, Hao Zhou, Jun Sun
HIGHLIGHT: In this paper, we propose W^3Net, which attempts to address above challenges by decomposing the pedestrian detection task into Where, What and Whether problem directing against pedestrian localization, scale prediction and classification correspondingly.

1398, TITLE: CoverNet: Multimodal Behavior Prediction Using Trajectory Sets
http://openaccess.thecvf.com/content_CVPR_2020/html/Phan-Minh_CoverNet_Multimodal_Behavior_Prediction_Using_Trajectory_Sets_CVPR_2020_paper.html
AUTHORS: Tung Phan-Minh, Elena Corina Grigore, Freddy A. Boulton, Oscar Beijbom, Eric M. Wolff
HIGHLIGHT: We present CoverNet, a new method for multimodal, probabilistic trajectory prediction for urban driving.

1399, TITLE: Real-World Person Re-Identification via Degradation Invariance Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Huang_Real-World_Person_Re-Identification_via_Degradation_Invariance_Learning_CVPR_2020_paper.html
AUTHORS: Yukun Huang, Zheng-Jun Zha, Xueyang Fu, Richang Hong, Liang Li
HIGHLIGHT: In this paper, to solve the above problem, we propose a degradation invariance learning framework for real-world person Re-ID.

1400, TITLE: Defending and Harnessing the Bit-Flip Based Adversarial Weight Attack

http://openaccess.thecvf.com/content_CVPR_2020/html/He_Defending_and_Harnessing_the_Bit-Flip_Based_Adversarial_Weight_Attack_CVPR_2020_paper.html
AUTHORS: Zhezhi He, Adnan Siraj Rakin, Jingtao Li, Chaitali Chakrabarti, Deliang Fan
HIGHLIGHT: In this work, we conduct comprehensive investigations on BFA and propose to leverage binarization-aware training and its relaxation -- piece-wise clustering as simple and effective countermeasures to BFA.

1401, TITLE: Adversarial Latent Autoencoders
http://openaccess.thecvf.com/content_CVPR_2020/html/Pidhorskyi_Adversarial_Latent_Autoencoders_CVPR_2020_paper.html
AUTHORS: Stanislav Pidhorskyi, Donald A. Adjeroh, Gianfranco Doretto
HIGHLIGHT: We introduce an autoencoder that tackles these issues jointly, which we call Adversarial Latent Autoencoder (ALAE).

1402, TITLE: Adaptive Fractional Dilated Convolution Network for Image Aesthetics Assessment
http://openaccess.thecvf.com/content_CVPR_2020/html/Chen_Adaptive_Fractional_Dilated_Convolution_Network_for_Image_Aesthetics_Assessment_CVPR_2020_paper.html
AUTHORS: Qiuyu Chen, Wei Zhang, Ning Zhou, Peng Lei, Yi Xu, Yu Zheng, Jianping Fan
HIGHLIGHT: In this paper, an adaptive fractional dilated convolution (AFDC), which is aspect-ratio-embedded, composition-preserving and parameter-free, is developed to tackle this issue natively in convolutional kernel level.

1403, TITLE: Deep Generative Model for Robust Imbalance Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Wang_Deep_Generative_Model_for_Robust_Imbalance_Classification_CVPR_2020_paper.html
AUTHORS: Xinyue Wang, Yilin Lyu, Liping Jing
HIGHLIGHT: In this paper, a deep generative classifier is proposed to mitigate this issue via both data perturbation and model perturbation.

1404, TITLE: Learning Deep Network for Detecting 3D Object Keypoints and 6D Poses
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Learning_Deep_Network_for_Detecting_3D_Object_Keypoints_and_6D_CVPR_2020_paper.html
AUTHORS: Wanqing Zhao, Shaobo Zhang, Ziyu Guan, Wei Zhao, Jinye Peng, Jianping Fan
HIGHLIGHT: In this paper, we develop a keypoint-based 6D object pose detection method (and its deep network) called Object Keypoint based POSe Estimation (OK-POSE).

1405, TITLE: MetaIQA: Deep Meta-Learning for No-Reference Image Quality Assessment
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhu_MetaIQA_Deep_Meta-Learning_for_No-Reference_Image_Quality_Assessment_CVPR_2020_paper.html
AUTHORS: Hancheng Zhu, Leida Li, Jinjian Wu, Weisheng Dong, Guangming Shi
HIGHLIGHT: With this motivation, this paper presents a no-reference IQA metric based on deep meta-learning.

1406, TITLE: Sketchformer: Transformer-Based Representation for Sketched Structure
http://openaccess.thecvf.com/content_CVPR_2020/html/Ribeiro_Sketchformer_Transformer-Based_Representation_for_Sketched_Structure_CVPR_2020_paper.html
AUTHORS: Leo Sampaio Ferraz Ribeiro, Tu Bui, John Collomosse, Moacir Ponti
HIGHLIGHT: Sketchformer is a novel transformer-based representation for encoding free-hand sketches input in a vector form, i.e. as a sequence of strokes.

1407, TITLE: Cylindrical Convolutional Networks for Joint Object Detection and Viewpoint Estimation
http://openaccess.thecvf.com/content_CVPR_2020/html/Joung_Cylindrical_Convolutional_Networks_for_Joint_Object_Detection_and_Viewpoint_Estimation_CVPR_2020_paper.html
AUTHORS: Sunghun Joung, Seungryong Kim, Hanjae Kim, Minsu Kim, Ig-Jae Kim, Junghyun Cho, Kwanghoon Sohn
HIGHLIGHT: To overcome this limitation, we introduce a learnable module, cylindrical convolutional networks (CCNs), that exploit cylindrical representation of a convolutional kernel defined in the 3D space.

1408, TITLE: Learning a Unified Sample Weighting Network for Object Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Cai_Learning_a_Unified_Sample_Weighting_Network_for_Object_Detection_CVPR_2020_paper.html
AUTHORS: Qi Cai, Yingwei Pan, Yu Wang, Jingen Liu, Ting Yao, Tao Mei
HIGHLIGHT: To this end, we devise a general loss function to cover most region-based object detectors with various sampling strategies, and then based on it we propose a unified sample weighting network to predict a sample's task weights.

1409, TITLE: Old Is Gold: Redefining the Adversarially Learned One-Class Classifier Training Paradigm

http://openaccess.thecvf.com/content_CVPR_2020/html/Zaheer_Old_Is_Gold_Redefining_the_Adversarially_Learned_One-Class_Classifier_Training_CVPR_2020_paper.html
AUTHORS: Muhammad Zaigham Zaheer, Jin-Ha Lee, Marcella Astrid, Seung-Ik Lee
HIGHLIGHT: In this study, we propose a framework that effectively generates stable results across a wide range of training steps and allows us to use both the generator and the discriminator of an adversarial model for efficient and robust anomaly detection.

1410, TITLE: An Adaptive Neural Network for Unsupervised Mosaic Consistency Analysis in Image Forensics
http://openaccess.thecvf.com/content_CVPR_2020/html/Bammey_An_Adaptive_Neural_Network_for_Unsupervised_Mosaic_Consistency_Analysis_in_CVPR_2020_paper.html
AUTHORS: Quentin Bammey, Rafael Grompone von Gioi, Jean-Michel Morel
HIGHLIGHT: In this paper we develop a blind method that can train directly on unlabelled and potentially forged images to point out local mosaic inconsistencies.

1411, TITLE: McFlow: Monte Carlo Flow Models for Data Imputation
http://openaccess.thecvf.com/content_CVPR_2020/html/Richardson_McFlow_Monte_Carlo_Flow_Models_for_Data_Imputation_CVPR_2020_paper.html
AUTHORS: Trevor W. Richardson, Wencheng Wu, Lei Lin, Beilei Xu, Edgar A. Bernal
HIGHLIGHT: To that end, we propose MCFlow, a deep framework for imputation that leverages normalizing flow generative models and Monte Carlo sampling.

1412, TITLE: Learning to See Through Obstructions
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Learning_to_See_Through_Obstructions_CVPR_2020_paper.html
AUTHORS: Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, Jia-Bin Huang
HIGHLIGHT: We present a learning-based approach for removing unwanted obstructions, such as window reflections, fence occlusions or raindrops, from a short sequence of images captured by a moving camera.

1413, TITLE: GaitPart: Temporal Part-Based Model for Gait Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Fan_GaitPart_Temporal_Part-Based_Model_for_Gait_Recognition_CVPR_2020_paper.html
AUTHORS: Chao Fan, Yunjie Peng, Chunshui Cao, Xu Liu, Saihui Hou, Jiannan Chi, Yongzhen Huang, Qing Li, Zhiqiang He
HIGHLIGHT: Then, we propose a novel part-based model GaitPart and get two aspects effect of boosting the performance: On the one hand, Focal Convolution Layer, a new applying of convolution, is presented to enhance the fine-grained learning of the part-level spatial features. On the other hand, the Micro-motion Capture Module (MCM) is proposed and there are several parallel MCMs in the GaitPart corresponding to the pre-defined parts of the human body, respectively.

1414, TITLE: EmotiCon: Context-Aware Multimodal Emotion Recognition Using Frege's Principle
http://openaccess.thecvf.com/content_CVPR_2020/html/Mittal_EmotiCon_Context-Aware_Multimodal_Emotion_Recognition_Using_Freges_Principle_CVPR_2020_paper.html
AUTHORS: Trisha Mittal, Pooja Guhan, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, Dinesh Manocha
HIGHLIGHT: We present EmotiCon, a learning-based algorithm for context-aware perceived human emotion recognition from videos and images.
We also introduce a new dataset, GroupWalk, which is a collection of videos captured in multiple real-world settings of people walking.

1415, TITLE: Can Deep Learning Recognize Subtle Human Activities?
http://openaccess.thecvf.com/content_CVPR_2020/html/Jacquot_Can_Deep_Learning_Recognize_Subtle_Human_Activities_CVPR_2020_paper.html
AUTHORS: Vincent Jacquot, Zhuofan Ying, Gabriel Kreiman
HIGHLIGHT: In this work, we propose a new action classification challenge that is performed well by humans, but poorly by state-of-the-art Deep Learning models.

1416, TITLE: PhysGAN: Generating Physical-World-Resilient Adversarial Examples for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2020/html/Kong_PhysGAN_Generating_Physical-World-Resilient_Adversarial_Examples_for_Autonomous_Driving_CVPR_2020_paper.html
AUTHORS: Zelun Kong, Junfeng Guo, Ang Li, Cong Liu
HIGHLIGHT: We present PhysGAN, which generates physical-world-resilient adversarial examples for misleading autonomous driving systems in a continuous manner.

1417, TITLE: ILFO: Adversarial Attack on Adaptive Neural Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Haque_ILFO_Adversarial_Attack_on_Adaptive_Neural_Networks_CVPR_2020_paper.html

AUTHORS: Mirazul Haque, Anki Chauhan, Cong Liu, Wei Yang
HIGHLIGHT: In this paper, we investigate the robustness of neural networks against energy-oriented attacks.

1418, TITLE: On Translation Invariance in CNNs: Convolutional Layers Can Exploit Absolute Spatial Location
http://openaccess.thecvf.com/content_CVPR_2020/html/Kayhan_On_Translation_Invariance_in_CNNs_Convolutional_Layers_Can_Exploit_Absolute_CVPR_2020_paper.html
AUTHORS: Osman Semih Kayhan, Jan C. van Gemert
HIGHLIGHT: In this paper we challenge the common assumption that convolutional layers in modern CNNs are translation invariant.

1419, TITLE: Diverse Image Generation via Self-Conditioned GANs
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_Diverse_Image_Generation_via_Self-Conditioned_GANs_CVPR_2020_paper.html
AUTHORS: Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, Antonio Torralba
HIGHLIGHT: We introduce a simple but effective unsupervised method for generating diverse images.

1420, TITLE: Inducing Hierarchical Compositional Model by Sparsifying Generator Network
http://openaccess.thecvf.com/content_CVPR_2020/html/Xing_Inducing_Hierarchical_Compositional_Model_by_Sparsifying_Generator_Network_CVPR_2020_paper.html
AUTHORS: Xianglei Xing, Tianfu Wu, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT: This paper proposes to learn hierarchical compositional AND-OR model for interpretable image synthesis by sparsifying the generator network.

1421, TITLE: CARP: Compression Through Adaptive Recursive Partitioning for Multi-Dimensional Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Liu_CARP_Compression_Through_Adaptive_Recursive_Partitioning_for_Multi-Dimensional_Images_CVPR_2020_paper.html
AUTHORS: Rongjie Liu, Meng Li, Li Ma
HIGHLIGHT: We present such a method for multi-dimensional image compression called Compression via Adaptive Recursive Partitioning (CARP).

1422, TITLE: GrappaNet: Combining Parallel Imaging With Deep Learning for Multi-Coil MRI Reconstruction
http://openaccess.thecvf.com/content_CVPR_2020/html/Sriram_GrappaNet_Combining_Parallel_Imaging_With_Deep_Learning_for_Multi-Coil_MRI_CVPR_2020_paper.html
AUTHORS: Anuroop Sriram, Jure Zbontar, Tullie Murrell, C. Lawrence Zitnick, Aaron Defazio, Daniel K. Sodickson
HIGHLIGHT: In this paper, we present a novel method to integrate traditional parallel imaging methods into deep neural networks that is able to generate high quality reconstructions even for high acceleration factors.

1423, TITLE: Can Weight Sharing Outperform Random Architecture Search? An Investigation With TuNAS
http://openaccess.thecvf.com/content_CVPR_2020/html/Bender_Can_Weight_Sharing_Outperform_Random_Architecture_Search_An_Investigation_With_CVPR_2020_paper.html
AUTHORS: Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, Quoc V. Le
HIGHLIGHT: While the efficacies of both methods are problem-dependent, our experiments demonstrate that there are large, realistic tasks where efficient search methods can provide substantial gains over random search.

1424, TITLE: Context Aware Graph Convolution for Skeleton-Based Action Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Context_Aware_Graph_Convolution_for_Skeleton-Based_Action_Recognition_CVPR_2020_paper.html
AUTHORS: Xikun Zhang, Chang Xu, Dacheng Tao
HIGHLIGHT: In this paper, we propose a context aware graph convolutional network (CA-GCN).

1425, TITLE: Fast(er) Reconstruction of Shredded Text Documents via Self-Supervised Deep Asymmetric Metric Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Paixao_Faster_Reconstruction_of_Shredded_Text_Documents_via_Self-Supervised_Deep_Asymmetric_CVPR_2020_paper.html
AUTHORS: Thiago M. Paixao, Rodrigo F. Berriel, Maria C. S. Boeres, Alessandro L. Koerich, Claudine Badue, Alberto F. De Souza, Thiago Oliveira-Santos
HIGHLIGHT: This work proposes a scalable deep learning approach for measuring pairwise compatibility in which the number of inferences scales linearly (rather than quadratically) with the number of shreds.

1426, TITLE: Revisiting Pose-Normalization for Fine-Grained Few-Shot Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Tang_Revisiting_Pose-Normalization_for_Fine-Grained_Few-Shot_Recognition_CVPR_2020_paper.html

AUTHORS: Luming Tang, Davis Wertheimer, Bharath Hariharan
HIGHLIGHT: A solution is to use pose-normalized representations: first localize semantic parts in each image, and then describe images by characterizing the appearance of each part. While such representations are out of favor for fully supervised classification, we show that they are extremely effective for few-shot fine-grained classification.

1427, TITLE: RankMI: A Mutual Information Maximizing Ranking Loss
http://openaccess.thecvf.com/content_CVPR_2020/html/Kemertas_RankMI_A_Mutual_Information_Maximizing_Ranking_Loss_CVPR_2020_paper.html
AUTHORS: Mete Kemertas, Leila Pishdad, Konstantinos G. Derpanis, Afsaneh Fazly
HIGHLIGHT: We introduce an information-theoretic loss function, RankMI, and an associated training algorithm for deep representation learning for image retrieval.

1428, TITLE: Learning Memory-Guided Normality for Anomaly Detection
http://openaccess.thecvf.com/content_CVPR_2020/html/Park_Learning_Memory-Guided_Normality_for_Anomaly_Detection_CVPR_2020_paper.html
AUTHORS: Hyunjong Park, Jongyoun Noh, Bumsub Ham
HIGHLIGHT: To address this problem, we present an unsupervised learning approach to anomaly detection that considers the diversity of normal patterns explicitly, while lessening the representation capacity of CNNs.

1429, TITLE: Appearance Shock Grammar for Fast Medial Axis Extraction From Real Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Camaro_Appearance_Shock_Grammar_for_Fast_Medial_Axis_Extraction_From_Real_CVPR_2020_paper.html
AUTHORS: Charles-Olivier Dufresne Camaro, Morteza Rezanejad, Stavros Tsogkas, Kaleem Siddiqi, Sven Dickinson
HIGHLIGHT: We combine ideas from shock graph theory with more recent appearance-based methods for medial axis extraction from complex natural scenes, improving upon the present best unsupervised method, in terms of efficiency and performance.

1430, TITLE: Generalizing Hand Segmentation in Egocentric Videos With Uncertainty-Guided Model Adaptation
http://openaccess.thecvf.com/content_CVPR_2020/html/Cai_Generalizing_Hand_Segmentation_in_Egocentric_Videos_With_Uncertainty-Guided_Model_Adaptation_CVPR_2020_paper.html
AUTHORS: Minjie Cai, Feng Lu, Yoichi Sato
HIGHLIGHT: In this work, we solve the hand segmentation generalization problem without requiring segmentation labels in the target domain.

1431, TITLE: DeFeat-Net: General Monocular Depth via Simultaneous Unsupervised Representation Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Spencer_DeFeat-Net_General_Monocular_Depth_via_Simultaneous_Unsupervised_Representation_Learning_CVPR_2020_paper.html
AUTHORS: Jaime Spencer, Richard Bowden, Simon Hadfield
HIGHLIGHT: We propose DeFeat-Net (Depth & Feature network), an approach to simultaneously learn a cross-domain dense feature representation, alongside a robust depth-estimation framework based on warped feature consistency.

1432, TITLE: Learning Visual Motion Segmentation Using Event Surfaces
http://openaccess.thecvf.com/content_CVPR_2020/html/Mitrokhin_Learning_Visual_Motion_Segmentation_Using_Event_Surfaces_CVPR_2020_paper.html
AUTHORS: Anton Mitrokhin, Zhiyuan Hua, Cornelia Fermuller, Yiannis Aloimonos
HIGHLIGHT: In this work we present a Graph Convolutional neural network for the task of scene motion segmentation by a moving camera.

1433, TITLE: Social-STGCNN: A Social Spatio-Temporal Graph Convolutional Neural Network for Human Trajectory Prediction
http://openaccess.thecvf.com/content_CVPR_2020/html/Mohamed_Social-STGCNN_A_Social_Spatio-Temporal_Graph_Convolutional_Neural_Network_for_Human_CVPR_2020_paper.html
AUTHORS: Abduallah Mohamed, Kun Qian, Mohamed Elhoseiny, Christian Claudel
HIGHLIGHT: We propose the Social Spatio-Temporal Graph Convolutional Neural Network (Social-STGCNN), which substitutes the need of aggregation methods by modeling the interactions as a graph.

1434, TITLE: Discriminative Multi-Modality Speech Recognition
http://openaccess.thecvf.com/content_CVPR_2020/html/Xu_Discriminative_Multi-Modality_Speech_Recognition_CVPR_2020_paper.html
AUTHORS: Bo Xu, Cheng Lu, Yandong Guo, Jacob Wang
HIGHLIGHT: In this paper, we propose a two-stage speech recognition model.

1435, TITLE:	Clean-Label Backdoor Attacks on Video Recognition Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Clean-
Label_Backdoor_Attacks_on_Video_Recognition_Models_CVPR_2020_paper.html
AUTHORS:	Shihao Zhao, Xingjun Ma, Xiang Zheng, James Bailey, Jingjing Chen, Yu-Gang Jiang
HIGHLIGHT:	In this paper, we show that existing image backdoor attacks are far less effective on videos, and outline 4 strict conditions where existing attacks are likely to fail: 1) scenarios with more input dimensions (eg. videos), 2) scenarios with high resolution, 3) scenarios with a large number of classes and few examples per class (a "sparse dataset"), and 4) attacks with access to correct labels (eg. clean-label attacks).

1436, TITLE:	Detecting Adversarial Samples Using Influence Functions and Nearest Neighbors
http://openaccess.thecvf.com/content_CVPR_2020/html/Cohen_Detecting_Adversarial_Samples_Using_Influence_Functions_and_N
earest_Neighbors_CVPR_2020_paper.html
AUTHORS:	Gilad Cohen, Guillermo Sapiro, Raja Giryes
HIGHLIGHT:	In this work, we present a method for detecting such adversarial attacks, which is suitable for any pre-trained neural network classifier.

1437, TITLE:	Unsupervised Model Personalization While Preserving Privacy and Scalability: An Open Problem
http://openaccess.thecvf.com/content_CVPR_2020/html/De_Lange_Unsupervised_Model_Personalization_While_Preserving_Privac
y_and_Scalability_An_Open_CVPR_2020_paper.html
AUTHORS:	Matthias De Lange, Xu Jia, Sarah Parisot, Ales Leonardis, Gregory Slabaugh, Tinne Tuytelaars
HIGHLIGHT:	We aim to address this challenge within the continual learning paradigm and provide a novel Dual User-Adaptation framework (DUA) to explore the problem.

1438, TITLE:	GIFnets: Differentiable GIF Encoding Framework
http://openaccess.thecvf.com/content_CVPR_2020/html/Yoo_GIFnets_Differentiable_GIF_Encoding_Framework_CVPR_2020_pape
r.html
AUTHORS:	Innfarn Yoo, Xiyang Luo, Yilin Wang, Feng Yang, Peyman Milanfar
HIGHLIGHT:	To reduce artifacts and provide a better and more efficient GIF encoding, we introduce a differentiable GIF encoding pipeline, which includes three novel neural networks: PaletteNet, DitherNet, and BandingNet.

1439, TITLE:	Learning Invariant Representation for Unsupervised Image Restoration
http://openaccess.thecvf.com/content_CVPR_2020/html/Du_Learning_Invariant_Representation_for_Unsupervised_Image_Restorati
on_CVPR_2020_paper.html
AUTHORS:	Wenchao Du, Hu Chen, Hongyu Yang
HIGHLIGHT:	Instead, we propose an unsupervised learning method that explicitly learns invariant presentation from noisy data and reconstructs clear observations.

1440, TITLE:	Improved Few-Shot Visual Classification
http://openaccess.thecvf.com/content_CVPR_2020/html/Bateni_Improved_Few-Shot_Visual_Classification_CVPR_2020_paper.html
AUTHORS:	Peyman Bateni, Raghav Goyal, Vaden Masrani, Frank Wood, Leonid Sigal
HIGHLIGHT:	In this paper, we explore the hypothesis that a simple class-covariance-based distance metric, namely the Mahalanobis distance, adopted into a state of the art few-shot learning approach (CNAPS) can, in and of itself, lead to a significant performance improvement.

1441, TITLE:	Learning Weighted Submanifolds With Variational Autoencoders and Riemannian Variational Autoencoders
http://openaccess.thecvf.com/content_CVPR_2020/html/Miolane_Learning_Weighted_Submanifolds_With_Variational_Autoencoder
s_and_Riemannian_Variational_Autoencoders_CVPR_2020_paper.html
AUTHORS:	Nina Miolane, Susan Holmes
HIGHLIGHT:	In this paper, we are interested in variants to learn potentially highly curved submanifolds of manifold-valued data.

1442, TITLE:	Learning Geocentric Object Pose in Oblique Monocular Images
http://openaccess.thecvf.com/content_CVPR_2020/html/Christie_Learning_Geocentric_Object_Pose_in_Oblique_Monocular_Images
_CVPR_2020_paper.html
AUTHORS:	Gordon Christie, Rodrigo Rene Rai Munoz Abujder, Kevin Foster, Shea Hagstrom, Gregory D. Hager, Myron Z. Brown
HIGHLIGHT:	Inspired by recent work in monocular height above ground prediction and optical flow prediction from static images, we develop an encoding of geocentric pose to address this challenge and train a deep network to compute the representation densely, supervised by publicly available airborne lidar.

1443, TITLE:	Understanding Adversarial Examples From the Mutual Influence of Images and Perturbations

http://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Understanding_Adversarial_Examples_From_the_Mutual_Influence_of_Images_and_CVPR_2020_paper.html
AUTHORS: Chaoning Zhang, Philipp Benz, Tooba Imtiaz, In So Kweon
HIGHLIGHT: We propose to treat the DNN logits as a vector for feature representation, and exploit them to analyze the mutual influence of two independent inputs based on the Pearson correlation coefficient (PCC).

1444, TITLE: Your Local GAN: Designing Two Dimensional Local Attention Mechanisms for Generative Models
http://openaccess.thecvf.com/content_CVPR_2020/html/Daras_Your_Local_GAN_Designing_Two_Dimensional_Local_Attention_Mechanisms_for_CVPR_2020_paper.html
AUTHORS: Giannis Daras, Augustus Odena, Han Zhang, Alexandros G. Dimakis
HIGHLIGHT: We introduce a new local sparse attention layer that preserves two-dimensional geometry and locality.

1445, TITLE: MoreFusion: Multi-object Reasoning for 6D Pose Estimation from Volumetric Fusion
http://openaccess.thecvf.com/content_CVPR_2020/html/Wada_MoreFusion_Multi-object_Reasoning_for_6D_Pose_Estimation_from_Volumetric_Fusion_CVPR_2020_paper.html
AUTHORS: Kentaro Wada, Edgar Sucar, Stephen James, Daniel Lenton, Andrew J. Davison
HIGHLIGHT: We present a system which can estimate the accurate poses of multiple known objects in contact and occlusion from real-time, embodied multi-view vision.

1446, TITLE: HCNAF: Hyper-Conditioned Neural Autoregressive Flow and its Application for Probabilistic Occupancy Map Forecasting
http://openaccess.thecvf.com/content_CVPR_2020/html/Oh_HCNAF_Hyper-Conditioned_Neural_Autoregressive_Flow_and_its_Application_for_Probabilistic_CVPR_2020_paper.html
AUTHORS: Geunseob Oh, Jean-Sebastien Valois
HIGHLIGHT: We introduce Hyper-Conditioned Neural Autoregressive Flow (HCNAF); a powerful universal distribution approximator designed to model arbitrarily complex conditional probability density functions.

1447, TITLE: Detail-recovery Image Deraining via Context Aggregation Networks
http://openaccess.thecvf.com/content_CVPR_2020/html/Deng_Detail-recovery_Image_Deraining_via_Context_Aggregation_Networks_CVPR_2020_paper.html
AUTHORS: Sen Deng, Mingqiang Wei, Jun Wang, Yidan Feng, Luming Liang, Haoran Xie, Fu Lee Wang, Meng Wang
HIGHLIGHT: We propose an end-to-end detail-recovery image deraining network (termed a DRDNet) to solve the problem.

1448, TITLE: MCEN: Bridging Cross-Modal Gap between Cooking Recipes and Dish Images with Latent Variable Model
http://openaccess.thecvf.com/content_CVPR_2020/html/Fu_MCEN_Bridging_Cross-Modal_Gap_between_Cooking_Recipes_and_Dish_Images_CVPR_2020_paper.html
AUTHORS: Han Fu, Rui Wu, Chenghao Liu, Jianling Sun
HIGHLIGHT: In this paper, we focus on the task of cross-modal retrieval between food images and cooking recipes.

1449, TITLE: Hypergraph Attention Networks for Multimodal Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Kim_Hypergraph_Attention_Networks_for_Multimodal_Learning_CVPR_2020_paper.html
AUTHORS: Eun-Sol Kim, Woo Young Kang, Kyoung-Woon On, Yu-Jung Heo, Byoung-Tak Zhang
HIGHLIGHT: To resolve this problem, we propose Hypergraph Attention Networks (HANs), which define a common semantic space among the modalities with symbolic graphs and extract a joint representation of the modalities based on a co-attention map constructed in the semantic space.

1450, TITLE: Moving in the Right Direction: A Regularization for Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2020/html/Mohan_Moving_in_the_Right_Direction_A_Regularization_for_Deep_Metric_CVPR_2020_paper.html
AUTHORS: Deen Dayal Mohan, Nishant Sankaran, Dennis Fedorishin, Srirangaraj Setlur, Venu Govindaraju
HIGHLIGHT: In this work, we identify a shortcoming of existing loss formulations which fail to consider more optimal directions of pair displacements as another criterion for optimization.

1451, TITLE: Rethinking Depthwise Separable Convolutions: How Intra-Kernel Correlations Lead to Improved MobileNets
http://openaccess.thecvf.com/content_CVPR_2020/html/Haase_Rethinking_Depthwise_Separable_Convolutions_How_Intra-Kernel_Correlations_Lead_to_Improved_CVPR_2020_paper.html
AUTHORS: Daniel Haase, Manuel Amthor
HIGHLIGHT: We introduce blueprint separable convolutions (BSConv) as highly efficient building blocks for CNNs.

1452, TITLE: Seeing without Looking: Contextual Rescoring of Object Detections for AP Maximization

http://openaccess.thecvf.com/content_CVPR_2020/html/Pato_Seeing_without_Looking_Contextual_Rescoring_of_Object_Detections
_for_AP_CVPR_2020_paper.html
AUTHORS: Lourenco V. Pato, Renato Negrinho, Pedro M. Q. Aguiar
HIGHLIGHT: We propose to incorporate context in object detection by post-processing the output of an arbitrary detector to rescore the confidences of its detections.

1453, TITLE: End-to-End Adversarial-Attention Network for Multi-Modal Clustering
http://openaccess.thecvf.com/content_CVPR_2020/html/Zhou_End-to-End_Adversarial-Attention_Network_for_Multi-
Modal_Clustering_CVPR_2020_paper.html
AUTHORS: Runwu Zhou, Yi-Dong Shen
HIGHLIGHT: In this paper, we present an End-to-end Adversarial-attention network for Multi-modal Clustering (EAMC), where adversarial learning and attention mechanism are leveraged to align the latent feature distributions and quantify the importance of modalities respectively.

1454, TITLE: Fast Sparse ConvNets
http://openaccess.thecvf.com/content_CVPR_2020/html/Elsen_Fast_Sparse_ConvNets_CVPR_2020_paper.html
AUTHORS: Erich Elsen, Marat Dukhan, Trevor Gale, Karen Simonyan
HIGHLIGHT: In this work, we further expand the arsenal of efficient building blocks for neural network architectures; but instead of combining standard primitives (such as convolution), we advocate for the replacement of these dense primitives with their sparse counterparts.

1455, TITLE: Few Sample Knowledge Distillation for Efficient Network Compression
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Few_Sample_Knowledge_Distillation_for_Efficient_Network_Compressi
on_CVPR_2020_paper.html
AUTHORS: Tianhong Li, Jianguo Li, Zhuang Liu, Changshui Zhang
HIGHLIGHT: This paper proposes a novel solution for knowledge distillation from label-free few samples to realize both data efficiency and training/processing efficiency.

1456, TITLE: Predicting Sharp and Accurate Occlusion Boundaries in Monocular Depth Estimation Using Displacement
Fields
http://openaccess.thecvf.com/content_CVPR_2020/html/Ramamonjisoa_Predicting_Sharp_and_Accurate_Occlusion_Boundaries_in_
Monocular_Depth_Estimation_CVPR_2020_paper.html
AUTHORS: Michael Ramamonjisoa, Yuming Du, Vincent Lepetit
HIGHLIGHT: We instead learn to predict, given a depth map predicted by some reconstruction method, a 2D displacement field able to re-sample pixels around the occlusion boundaries into sharper reconstructions.

1457, TITLE: Shape correspondence using anisotropic Chebyshev spectral CNNs
http://openaccess.thecvf.com/content_CVPR_2020/html/Li_Shape_correspondence_using_anisotropic_Chebyshev_spectral_CNNs_C
VPR_2020_paper.html
AUTHORS: Qinsong Li, Shengjun Liu, Ling Hu, Xinru Liu
HIGHLIGHT: In this paper, we propose a novel architecture for shape correspondence, termed Anisotropic Chebyshev spectral CNNs (ACSCNNs), based on a new extension of the manifold convolution operator.

1458, TITLE: RetinaTrack: Online Single Stage Joint Detection and Tracking
http://openaccess.thecvf.com/content_CVPR_2020/html/Lu_RetinaTrack_Online_Single_Stage_Joint_Detection_and_Tracking_CVP
R_2020_paper.html
AUTHORS: Zhichao Lu, Vivek Rathod, Ronny Votel, Jonathan Huang
HIGHLIGHT: In this paper we focus on the tracking-by-detection paradigm for autonomous driving where both tasks are mission critical.

1459, TITLE: Multimodal Categorization of Crisis Events in Social Media
http://openaccess.thecvf.com/content_CVPR_2020/html/Abavisani_Multimodal_Categorization_of_Crisis_Events_in_Social_Media_
CVPR_2020_paper.html
AUTHORS: Mahdi Abavisani, Liwei Wu, Shengli Hu, Joel Tetreault, Alejandro Jaimes
HIGHLIGHT: In this paper, we present a new multimodal fusion method that leverages both images and texts as input.

1460, TITLE: SPARE3D: A Dataset for SPAtial REasoning on Three-View Line Drawings
http://openaccess.thecvf.com/content_CVPR_2020/html/Han_SPARE3D_A_Dataset_for_SPAtial_REasoning_on_Three-
View_Line_Drawings_CVPR_2020_paper.html
AUTHORS: Wenyu Han, Siyuan Xiang, Chenhui Liu, Ruoyu Wang, Chen Feng
HIGHLIGHT: Can deep networks be trained to perform spatial reasoning tasks? How can we measure their "spatial intelligence"? To answer these questions, we present the SPARE3D dataset.

1461, TITLE: SwapText: Image Based Texts Transfer in Scenes
http://openaccess.thecvf.com/content_CVPR_2020/html/Yang_SwapText_Image_Based_Texts_Transfer_in_Scenes_CVPR_2020_paper.html
AUTHORS: Qiangpeng Yang, Jun Huang, Wei Lin
HIGHLIGHT: In this work, we present SwapText, a three-stage framework to transfer texts across scene images.

1462, TITLE: OrigamiNet: Weakly-Supervised, Segmentation-Free, One-Step, Full Page Text Recognition by learning to unfold
http://openaccess.thecvf.com/content_CVPR_2020/html/Yousef_OrigamiNet_Weakly-Supervised_Segmentation-Free_One-Step_Full_Page_Text_Recognition_by_learning_CVPR_2020_paper.html
AUTHORS: Mohamed Yousef, Tom E. Bishop
HIGHLIGHT: We propose a novel and simple neural network module, termed OrigamiNet, that can augment any CTC-trained, fully convolutional single line text recognizer, to convert it into a multi-line version by providing the model with enough spatial capacity to be able to properly collapse a 2D input signal into 1D without losing information.

1463, TITLE: FroDO: From Detections to 3D Objects
http://openaccess.thecvf.com/content_CVPR_2020/html/Runz_FroDO_From_Detections_to_3D_Objects_CVPR_2020_paper.html
AUTHORS: Martin Runz, Kejie Li, Meng Tang, Lingni Ma, Chen Kong, Tanner Schmidt, Ian Reid, Lourdes Agapito, Julian Straub, Steven Lovegrove, Richard Newcombe
HIGHLIGHT: We introduce FroDO, a method for accurate 3D reconstruction of object instances from RGB video that infers their location, pose and shape in a coarse to fine manner.

1464, TITLE: Single-Step Adversarial Training With Dropout Scheduling
http://openaccess.thecvf.com/content_CVPR_2020/html/B.S._Single-Step_Adversarial_Training_With_Dropout_Scheduling_CVPR_2020_paper.html
AUTHORS: Vivek B.S., R. Venkatesh Babu
HIGHLIGHT: In this work, (i) we show that models trained using single-step adversarial training method learn to prevent the generation of single-step adversaries, and this is due to over-fitting of the model during the initial stages of training, and (ii) to mitigate this effect, we propose a single-step adversarial training method with dropout scheduling.

1465, TITLE: Learning to Super Resolve Intensity Images From Events
http://openaccess.thecvf.com/content_CVPR_2020/papers/I._Learning_to_Super_Resolve_Intensity_Images_From_Events_CVPR_2020_paper.pdf
AUTHORS: S. Mohammad Mostafavi I., Jonghyun Choi, Kuk-Jin Yoon
HIGHLIGHT: We propose an end-to-end network to reconstruct high resolution, high dynamic range (HDR) images directly from the event stream.

1466, TITLE: DeeperForensics-1.0: A Large-Scale Dataset for Real-World Face Forgery Detection
http://openaccess.thecvf.com/content_CVPR_2020/papers/Jiang_DeeperForensics-1.0_A_Large-Scale_Dataset_for_Real-World_Face_Forgery_Detection_CVPR_2020_paper.pdf
AUTHORS: Liming Jiang, Ren Li, Wayne Wu, Chen Qian, Chen Change Loy
HIGHLIGHT: In this paper, we present our on-going effort of constructing a large-scale benchmark, DeeperForensics-1.0, for face forgery detection.

1467, TITLE: CNN-Generated Images Are Surprisingly Easy to Spot... for Now
http://openaccess.thecvf.com/content_CVPR_2020/papers/Wang_CNN-Generated_Images_Are_Surprisingly_Easy_to_Spot..._for_Now_CVPR_2020_paper.pdf
AUTHORS: Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, Alexei A. Efros
HIGHLIGHT: In this work we ask whether it is possible to create a ``universal'' detector for telling apart real images from these generated by a CNN, regardless of architecture or dataset used.