

- 1, TITLE: Embodied Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Das_Embodied_Question_Answering_CVPR_2018_paper.html
AUTHORS: Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra
HIGHLIGHT: In this work, we develop a dataset of questions and answers in House3D environments, evaluation metrics, and a hierarchical model trained with imitation and reinforcement learning.

- 2, TITLE: Learning by Asking Questions
http://openaccess.thecvf.com/content_cvpr_2018/html/Misra_Learning_by_Asking_CVPR_2018_paper.html
AUTHORS: Ishan Misra, Ross Girshick, Rob Fergus, Martial Hebert, Abhinav Gupta, Laurens van der Maaten
HIGHLIGHT: We introduce an interactive learning framework for the development and testing of intelligent visual systems, called learning-by-asking (LBA).

- 3, TITLE: Finding Tiny Faces in the Wild With Generative Adversarial Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Bai_Finding_Tiny_Faces_CVPR_2018_paper.html
AUTHORS: Yancheng Bai, Yongqiang Zhang, Mingli Ding, Bernard Ghanem
HIGHLIGHT: In this paper, we proposed an algorithm to directly generate a clear high-resolution face from a blurry small one by adopting a generative adversarial network (GAN).

- 4, TITLE: Learning Face Age Progression: A Pyramid Architecture of GANs
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Learning_Face_Age_CVPR_2018_paper.html
AUTHORS: Hongyu Yang, Di Huang, Yunhong Wang, Anil K. Jain
HIGHLIGHT: In this paper, we present a novel generative adversarial network based approach.

- 5, TITLE: PairedCycleGAN: Asymmetric Style Transfer for Applying and Removing Makeup
http://openaccess.thecvf.com/content_cvpr_2018/html/Chang_PairedCycleGAN_Asymmetric_Style_CVPR_2018_paper.html
AUTHORS: Huiwen Chang, Jingwan Lu, Fisher Yu, Adam Finkelstein
HIGHLIGHT: This paper introduces an automatic method for editing a portrait photo so that the subject appears to be wearing makeup in the style of another person in a reference photo.

- 6, TITLE: GANerated Hands for Real-Time 3D Hand Tracking From Monocular RGB
http://openaccess.thecvf.com/content_cvpr_2018/html/Mueller_GANerated_Hands_for_CVPR_2018_paper.html
AUTHORS: Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, Christian Theobalt
HIGHLIGHT: For training our CNN we propose a novel approach for the synthetic generation of training data that is based on a geometrically consistent image-to-image translation network.

- 7, TITLE: Learning Pose Specific Representations by Predicting Different Views
http://openaccess.thecvf.com/content_cvpr_2018/html/Poier_Learning_Pose_Specific_CVPR_2018_paper.html
AUTHORS: Georg Poier, David Schinagl, Horst Bischof
HIGHLIGHT: To address this issue, we develop a method to learn representations, which are very specific for articulated poses, without the need for labeled training data.

- 8, TITLE: Weakly and Semi Supervised Human Body Part Parsing via Pose-Guided Knowledge Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Fang_Weakly_and_Semi_CVPR_2018_paper.html
AUTHORS: Hao-Shu Fang, Guansong Lu, Xiaolin Fang, Jianwen Xie, Yu-Wing Tai, Cewu Lu
HIGHLIGHT: In this paper, we present a novel method to generate synthetic human part segmentation data using easily-obtained human keypoint annotations.

- 9, TITLE: Person Transfer GAN to Bridge Domain Gap for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Person_Transfer_GAN_CVPR_2018_paper.html
AUTHORS: Longhui Wei, Shiliang Zhang, Wen Gao, Qi Tian
HIGHLIGHT: To relieve the expensive costs of annotating new training samples, we propose a Person Transfer Generative Adversarial Network (PTGAN) to bridge the domain gap.

- 10, TITLE: Cross-Modal Deep Variational Hand Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Spurr_Cross-Modal_Deep_Variational_CVPR_2018_paper.html
AUTHORS: Adrian Spurr, Jie Song, Seonwook Park, Otmar Hilliges
HIGHLIGHT: In this work we propose a method to learn a statistical hand model represented by a cross-modal trained latent space via a generative deep neural network.

- 11, TITLE: Disentangled Person Image Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Ma_Disentangled_Person_Image_CVPR_2018_paper.html
AUTHORS: Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, Mario Fritz
HIGHLIGHT: In this work, we aim at generating such images based on a novel, two-stage reconstruction pipeline that learns a disentangled representation of the aforementioned image factors and generates novel person images at the same time.
- 12, TITLE: Super-FAN: Integrated Facial Landmark Localization and Super-Resolution of Real-World Low Resolution Faces in Arbitrary Poses With GANs
http://openaccess.thecvf.com/content_cvpr_2018/html/Bulat_Super-FAN_Integrated_Facial_CVPR_2018_paper.html
AUTHORS: Adrian Bulat, Georgios Tzimiropoulos
HIGHLIGHT: To this end, we make the following 5 contributions: (a) we propose Super-FAN: the very first end-to-end system that addresses both tasks simultaneously, i.e. both improves face resolution and detects the facial landmarks.
- 13, TITLE: Multistage Adversarial Losses for Pose-Based Human Image Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Si_Multistage_Adversarial_Losses_CVPR_2018_paper.html
AUTHORS: Chenyang Si, Wei Wang, Liang Wang, Tieniu Tan
HIGHLIGHT: In this paper, we propose a pose-based human image synthesis method which can keep the human posture unchanged in novel viewpoints.
- 14, TITLE: Rotation Averaging and Strong Duality
http://openaccess.thecvf.com/content_cvpr_2018/html/Eriksson_Rotation_Averaging_and_CVPR_2018_paper.html
AUTHORS: Anders Eriksson, Carl Olsson, Fredrik Kahl, Tat-Jun Chin
HIGHLIGHT: In this paper we explore the role of duality principles within the problem of rotation averaging, a fundamental task in a wide range of computer vision applications.
- 15, TITLE: Hybrid Camera Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Camposco_Hybrid_Camera_Pose_CVPR_2018_paper.html
AUTHORS: Federico Camposco, Andrea Cohen, Marc Pollefeys, Torsten Sattler
HIGHLIGHT: In this paper, we aim to solve the pose estimation problem of calibrated pinhole and generalized cameras w.r.t. a Structure-from-Motion (SfM) model by leveraging both 2D-3D correspondences as well as 2D-2D correspondences.
- 16, TITLE: A Certifiably Globally Optimal Solution to the Non-Minimal Relative Pose Problem
http://openaccess.thecvf.com/content_cvpr_2018/html/Briales_A_Certifiably_Globally_CVPR_2018_paper.html
AUTHORS: Jesus Briales, Laurent Kneip, Javier Gonzalez-Jimenez
HIGHLIGHT: This, notably, is the contribution of the present paper.
- 17, TITLE: Single View Stereo Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Luo_Single_View_Stereo_CVPR_2018_paper.html
AUTHORS: Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, Liang Lin
HIGHLIGHT: In this paper, we show for the first time that the monocular depth estimation problem can be reformulated as two sub-problems, a view synthesis procedure followed by stereo matching, with two intriguing properties, namely i) geometrical constraints can be explicitly imposed during inference; ii) demand on labelled depth data can be greatly alleviated.
- 18, TITLE: Fight Ill-Posedness With Ill-Posedness: Single-Shot Variational Depth Super-Resolution From Shading
http://openaccess.thecvf.com/content_cvpr_2018/html/Haefner_Fight_Ill-Posedness_With_CVPR_2018_paper.html
AUTHORS: Bjoern Haefner, Yvain Quéau, Thomas Möllenhoff, Daniel Cremers
HIGHLIGHT: We combine heterogeneous depth and color data in order to jointly solve the ill-posed depth super-resolution and shape-from-shading problems.
- 19, TITLE: Deep Depth Completion of a Single RGB-D Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Deep_Depth_Completion_CVPR_2018_paper.html
AUTHORS: Yinda Zhang, Thomas Funkhouser
HIGHLIGHT: The goal of our work is to complete the depth channel of an RGB-D image.
- 20, TITLE: Multi-View Harmonized Bilinear Network for 3D Object Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Multi-View_Harmonized_Bilinear_CVPR_2018_paper.html
AUTHORS: Tan Yu, Jingjing Meng, Junsong Yuan
HIGHLIGHT: Different from existing view-based methods pooling the view-wise features, we tackle this problem from the perspective of patches-to-patches similarity measurement.

- 21, TITLE: PPFNet: Global Context Aware Local Features for Robust 3D Point Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Deng_PPFNet_Global_Context_CVPR_2018_paper.html
AUTHORS: Haowen Deng, Tolga Birdal, Slobodan Ilic
HIGHLIGHT: We present PPFNet - Point Pair Feature NETWORK for deeply learning a globally informed 3D local feature descriptor to find correspondences in unorganized point clouds.
- 22, TITLE: FoldingNet: Point Cloud Auto-Encoder via Deep Grid Deformation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_FoldingNet_Point_Cloud_CVPR_2018_paper.html
AUTHORS: Yaoqing Yang, Chen Feng, Yiru Shen, Dong Tian
HIGHLIGHT: In this work, a novel end-to-end deep auto-encoder is proposed to address unsupervised learning challenges on point clouds.
- 23, TITLE: A Papier-Mâché Approach to Learning 3D Surface Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Groueix_A_Papier-Mache_Approach_CVPR_2018_paper.html
AUTHORS: Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, Mathieu Aubry
HIGHLIGHT: We introduce a method for learning to generate the surface of 3D shapes.
- 24, TITLE: LEGO: Learning Edge With Geometry All at Once by Watching Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_LEGO_Learning_Edge_CVPR_2018_paper.html
AUTHORS: Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia
HIGHLIGHT: In this paper, we introduce a “3D as-smooth-as-possible (3D-ASAP)” prior inside the pipeline, which enables joint estimation of edges and 3D scene, yielding results with significant improvement in accuracy for fine detailed structures.
- 25, TITLE: Five-Point Fundamental Matrix Estimation for Uncalibrated Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Barath_Five-Point_Fundamental_Matrix_CVPR_2018_paper.html
AUTHORS: Daniel Barath
HIGHLIGHT: We aim at estimating the fundamental matrix in two views from five correspondences of rotation invariant features obtained by e.g. the SIFT detector.
- 26, TITLE: PointFusion: Deep Sensor Fusion for 3D Bounding Box Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_PointFusion_Deep_Sensor_CVPR_2018_paper.html
AUTHORS: Danfei Xu, Dragomir Anguelov, Ashesh Jain
HIGHLIGHT: We present PointFusion, a generic 3D object detection method that leverages both image and 3D point cloud information.
- 27, TITLE: Scalable Dense Non-Rigid Structure-From-Motion: A Grassmannian Perspective
http://openaccess.thecvf.com/content_cvpr_2018/html/Kumar_Scalable_Dense_Non-Rigid_CVPR_2018_paper.html
AUTHORS: Suryansh Kumar, Anoop Cherian, Yuchao Dai, Hongdong Li
HIGHLIGHT: To address these issues, in this paper, we propose a new approach for dense NRSfM by modeling the problem on a Grassmann manifold.
- 28, TITLE: GVCNN: Group-View Convolutional Neural Networks for 3D Shape Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Feng_GVCNN_Group-View_Convolutional_CVPR_2018_paper.html
AUTHORS: Yifan Feng, Zizhao Zhang, Xibin Zhao, Rongrong Ji, Yue Gao
HIGHLIGHT: To tackle this issue, in this paper, we propose a group-view convolutional neural network (GVCNN) framework for hierarchical correlation modeling towards discriminative 3D shape description.
- 29, TITLE: Depth and Transient Imaging With Compressive SPAD Array Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Depth_and_Transient_CVPR_2018_paper.html
AUTHORS: Qilin Sun, Xiong Dun, Yifan Peng, Wolfgang Heidrich
HIGHLIGHT: In this work we aim to overcome the spatial resolution limit of SPAD arrays by employing a compressive sensing camera design.
- 30, TITLE: GeoNet: Geometric Neural Network for Joint Depth and Surface Normal Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_GeoNet_Geometric_Neural_CVPR_2018_paper.html
AUTHORS: Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, Jiaya Jia
HIGHLIGHT: In this paper, we propose Geometric Neural Network (GeoNet) to jointly predict depth and surface normal maps from a single image.
- 31, TITLE: Real-Time Seamless Single Shot 6D Object Pose Prediction

- http://openaccess.thecvf.com/content_cvpr_2018/html/Tekin_Real-Time_Seamless_Single_CVPR_2018_paper.html
AUTHORS: Bugra Tekin, Sudipta N. Sinha, Pascal Fua
HIGHLIGHT: We propose a single-shot approach for simultaneously detecting an object in an RGB image and predicting its 6D pose without requiring multiple stages or having to examine multiple hypotheses.
- 32, TITLE: Factoring Shape, Pose, and Layout From the 2D Image of a 3D Scene
http://openaccess.thecvf.com/content_cvpr_2018/html/Tulsiani_Factoring_Shape_Pose_CVPR_2018_paper.html
AUTHORS: Shubham Tulsiani, Saurabh Gupta, David F. Fouhey, Alexei A. Efros, Jitendra Malik
HIGHLIGHT: The goal of this paper is to take a single 2D image of a scene and recover the 3D structure in terms of a small set of factors: a layout representing the enclosing surfaces as well as a set of objects represented in terms of shape and pose.
- 33, TITLE: Monocular Relative Depth Perception With Web Stereo Data Supervision
http://openaccess.thecvf.com/content_cvpr_2018/html/Xian_Monocular_Relative_Depth_CVPR_2018_paper.html
AUTHORS: Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, Zhenbo Luo
HIGHLIGHT: In this paper we study the problem of monocular relative depth perception in the wild. We introduce a simple yet effective method to automatically generate dense relative depth annotations from web stereo images, and propose a new dataset that consists of diverse images as well as corresponding dense relative depth maps.
- 34, TITLE: Spline Error Weighting for Robust Visual-Inertial Fusion
http://openaccess.thecvf.com/content_cvpr_2018/html/Ovren_Spline_Error_Weighting_CVPR_2018_paper.html
AUTHORS: Hannes Ovrén, Per-Erik Forssén
HIGHLIGHT: In this paper we derive and test a probability-based weighting that can balance residuals of different types in spline fitting.
- 35, TITLE: Single-Image Depth Estimation Based on Fourier Domain Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_Single-Image_Depth_Estimation_CVPR_2018_paper.html
AUTHORS: Jae-Han Lee, Minhyeok Heo, Kyung-Rae Kim, Chang-Su Kim
HIGHLIGHT: We propose a deep learning algorithm for single-image depth estimation based on the Fourier frequency domain analysis.
- 36, TITLE: Unsupervised Learning of Monocular Depth Estimation and Visual Odometry With Deep Feature Reconstruction
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhan_Unsupervised_Learning_of_CVPR_2018_paper.html
AUTHORS: Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, Ian Reid
HIGHLIGHT: In this paper, we explore the use of stereo sequences for learning depth and visual odometry.
- 37, TITLE: Detect-and-Track: Efficient Pose Estimation in Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Girdhar_Detect-and-Track_Efficient_Pose_CVPR_2018_paper.html
AUTHORS: Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, Du Tran
HIGHLIGHT: We propose an extremely lightweight yet highly effective approach that builds upon the latest advancements in human detection and video understanding.
- 38, TITLE: Supervision-by-Registration: An Unsupervised Approach to Improve the Precision of Facial Landmark Detectors
http://openaccess.thecvf.com/content_cvpr_2018/html/Dong_Supervision-by-Registration_An_Unsupervised_CVPR_2018_paper.html
AUTHORS: Xuanyi Dong, Shoou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, Yaser Sheikh
HIGHLIGHT: In this paper, we present supervision-by-registration, an unsupervised approach to improve the precision of facial landmark detectors on both images and video.
- 39, TITLE: Diversity Regularized Spatiotemporal Attention for Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Diversity_Regularized_Spatiotemporal_CVPR_2018_paper.html
AUTHORS: Shuang Li, Slawomir Bak, Peter Carr, Xiaogang Wang
HIGHLIGHT: Instead, we propose a new spatiotemporal attention model that automatically discovers a diverse set of distinctive body parts.
- 40, TITLE: Style Aggregated Network for Facial Landmark Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Dong_Style_Aggregated_Network_CVPR_2018_paper.html
AUTHORS: Xuanyi Dong, Yan Yan, Wanli Ouyang, Yi Yang
HIGHLIGHT: In this work, we propose a style-aggregated approach to deal with the large intrinsic variance of image styles for facial landmark detection.

- 41, TITLE: Learning Deep Models for Face Anti-Spoofing: Binary or Auxiliary Supervision
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Learning_Deep_Models_CVPR_2018_paper.html
AUTHORS: Yaojie Liu, Amin Jourabloo, Xiaoming Liu
HIGHLIGHT: In this paper, we argue the importance of auxiliary supervision to guide the learning toward discriminative and generalizable cues.
Further, we introduce a new face anti-spoofing database that covers a large range of illumination, subject, and pose variations.
- 42, TITLE: Deep Cost-Sensitive and Order-Preserving Feature Learning for Cross-Population Age Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Deep_Cost-Sensitive_and_CVPR_2018_paper.html
AUTHORS: Kai Li, Junliang Xing, Chi Su, Weiming Hu, Yundong Zhang, Stephen Maybank
HIGHLIGHT: In this work, we propose a Deep Cross-Population (DCP) age estimation model to achieve this goal.
- 43, TITLE: First-Person Hand Action Benchmark With RGB-D Videos and 3D Hand Pose Annotations
http://openaccess.thecvf.com/content_cvpr_2018/html/Garcia-Hernando_First-Person_Hand_Action_CVPR_2018_paper.html
AUTHORS: Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, Tae-Kyun Kim
HIGHLIGHT: In this work we study the use of 3D hand poses to recognize first-person dynamic hand actions interacting with 3D objects.
- 44, TITLE: A Pose-Sensitive Embedding for Person Re-Identification With Expanded Cross Neighborhood Re-Ranking
http://openaccess.thecvf.com/content_cvpr_2018/html/Sarfraz_A_Pose-Sensitive_Embedding_CVPR_2018_paper.html
AUTHORS: M. Saquib Sarfraz, Arne Schumann, Andreas Eberle, Rainer Stiefelhagen
HIGHLIGHT: In this paper we propose an effective approach that incorporates both the fine and coarse pose information of the person to learn a discriminative embedding.
- 45, TITLE: Disentangling 3D Pose in a Dendritic CNN for Unconstrained 2D Face Alignment
http://openaccess.thecvf.com/content_cvpr_2018/html/Kumar_Disentangling_3D_Pose_CVPR_2018_paper.html
AUTHORS: Amit Kumar, Rama Chellappa
HIGHLIGHT: In this paper, we present a single dendritic CNN, termed as Pose Conditioned Dendritic Convolution Neural Network (PCD-CNN), where a classification network is followed by a second and modular classification network, trained in an end to end fashion to obtain accurate landmark points.
- 46, TITLE: A Hierarchical Generative Model for Eye Image Synthesis and Eye Gaze Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_A_Hierarchical_Generative_CVPR_2018_paper.html
AUTHORS: Kang Wang, Rui Zhao, Qiang Ji
HIGHLIGHT: In this work, we introduce a Hierarchical Generative Model (HGM) to enable realistic forward eye image synthesis, as well as effective backward eye gaze estimation.
- 47, TITLE: MiCT: Mixed 3D/2D Convolutional Tube for Human Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_MiCT_Mixed_3D2D_CVPR_2018_paper.html
AUTHORS: Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, Wenjun Zeng
HIGHLIGHT: We thus propose a Mixed Convolutional Tube (MiCT) that integrates 2D CNNs with the 3D convolution module to generate deeper and more informative feature maps, while reducing training complexity in each round of spatio-temporal fusion.
- 48, TITLE: Learning to Estimate 3D Human Pose and Shape From a Single Color Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Pavlakos_Learning_to_Estimate_CVPR_2018_paper.html
AUTHORS: Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, Kostas Daniilidis
HIGHLIGHT: Our work aims to bridge this gap and proposes an efficient and effective direct prediction method based on ConvNets.
- 49, TITLE: Glimpse Clouds: Human Activity Recognition From Unstructured Feature Points
http://openaccess.thecvf.com/content_cvpr_2018/html/Baradel_Glimpse_Clouds_Human_CVPR_2018_paper.html
AUTHORS: Fabien Baradel, Christian Wolf, Julien Mille, Graham W. Taylor
HIGHLIGHT: We propose a method for human activity recognition from RGB data that does not rely on any pose information during test time, and does not explicitly calculate pose information internally.
- 50, TITLE: Context-Aware Deep Feature Compression for High-Speed Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Choi_Context-Aware_Deep_Feature_CVPR_2018_paper.html

AUTHORS: Jongwon Choi, Hyung Jin Chang, Tobias Fischer, Sangdoon Yun, Kyuewang Lee, Jiyeoup Jeong, Yiannis Demiris, Jin Young Choi

HIGHLIGHT: We propose a new context-aware correlation filter based tracking framework to achieve both high computational speed and state-of-the-art performance among real-time trackers.

51, **TITLE:** Correlation Tracking via Joint Discrimination and Reliability Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Correlation_Tracking_via_CVPR_2018_paper.html

AUTHORS: Chong Sun, Dong Wang, Huchuan Lu, Ming-Hsuan Yang

HIGHLIGHT: To address this issue, we propose a novel CF-based optimization problem to jointly model the discrimination and reliability information.

52, **TITLE:** PhaseNet for Video Frame Interpolation

http://openaccess.thecvf.com/content_cvpr_2018/html/Meyer_PhaseNet_for_Video_CVPR_2018_paper.html

AUTHORS: Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, Christopher Schroers

HIGHLIGHT: We propose a new approach, PhaseNet, that is designed to robustly handle challenging scenarios while also coping with larger motion.

53, **TITLE:** The Best of Both Worlds: Combining CNNs and Geometric Constraints for Hierarchical Motion Segmentation

http://openaccess.thecvf.com/content_cvpr_2018/html/Bideau_The_Best_of_CVPR_2018_paper.html

AUTHORS: Pia Bideau, Aruni RoyChowdhury, Rakesh R. Menon, Erik Learned-Miller

HIGHLIGHT: In this work, we build a new statistical model of rigid motion flow based on classical perspective projection constraints.

54, **TITLE:** Hyperparameter Optimization for Tracking With Continuous Deep Q-Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Dong_Hyperparameter_Optimization_for_CVPR_2018_paper.html

AUTHORS: Xingping Dong, Jianbing Shen, Wenguan Wang, Yu Liu, Ling Shao, Fatih Porikli

HIGHLIGHT: Here, we propose a novel hyperparameter optimization method that can find optimal hyperparameters for a given sequence using an action-prediction network leveraged on Continuous Deep Q-Learning.

55, **TITLE:** Scale-Transferrable Object Detection

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Scale-Transferrable_Object_Detection_CVPR_2018_paper.html

AUTHORS: Peng Zhou, Bingbing Ni, Cong Geng, Jianguo Hu, Yi Xu

HIGHLIGHT: In this work, we develop a novel Scale-Transferrable Detection Network (STDN) for detecting multi-scale objects in images.

56, **TITLE:** A Prior-Less Method for Multi-Face Tracking in Unconstrained Videos

http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_A_Prior-Less_Method_CVPR_2018_paper.html

AUTHORS: Chung-Ching Lin, Ying Hung

HIGHLIGHT: This paper presents a prior-less method for tracking and clustering an unknown number of human faces and maintaining their individual identities in unconstrained videos.

57, **TITLE:** End-to-End Flow Correlation Tracking With Spatial-Temporal Attention

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_End-to-End_Flow_Correlation_CVPR_2018_paper.html

AUTHORS: Zheng Zhu, Wei Wu, Wei Zou, Junjie Yan

HIGHLIGHT: In this paper, we propose the FlowTrack, which focuses on making use of the rich flow information in consecutive frames to improve the feature representation and the tracking accuracy.

58, **TITLE:** Deep Texture Manifold for Ground Terrain Recognition

http://openaccess.thecvf.com/content_cvpr_2018/html/Xue_Deep_Texture_Manifold_CVPR_2018_paper.html

AUTHORS: Jia Xue, Hang Zhang, Kristin Dana

HIGHLIGHT: We present a texture network called Deep Encoding Pooling Network (DEP) for the task of ground terrain recognition.

59, **TITLE:** Learning Superpixels With Segmentation-Aware Affinity Loss

http://openaccess.thecvf.com/content_cvpr_2018/html/Tu_Learning_Superpixels_With_CVPR_2018_paper.html

AUTHORS: Wei-Chih Tu, Ming-Yu Liu, Varun Jampani, Deqing Sun, Shao-Yi Chien, Ming-Hsuan Yang, Jan Kautz

HIGHLIGHT: In this work, we leverage deep neural networks to facilitate extracting superpixels from images.

60, **TITLE:** Interactive Image Segmentation With Latent Diversity

- http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Interactive_Image_Segmentation_CVPR_2018_paper.html
AUTHORS: Zhuwen Li, Qifeng Chen, Vladlen Koltun
HIGHLIGHT: We present an end-to-end learning approach to interactive image segmentation that tackles this ambiguity.
- 61, TITLE: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html
AUTHORS: Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, Oliver Wang
HIGHLIGHT: To answer these questions, we introduce a new dataset of human perceptual similarity judgments.
- 62, TITLE: Local Descriptors Optimized for Average Precision
http://openaccess.thecvf.com/content_cvpr_2018/html/He_Local_Descriptors_Optimized_CVPR_2018_paper.html
AUTHORS: Kun He, Yan Lu, Stan Sclaroff
HIGHLIGHT: In this paper, we improve the learning of local feature descriptors by optimizing the performance of descriptor matching, which is a common stage that follows descriptor extraction in local feature based pipelines, and can be formulated as nearest neighbor retrieval.
- 63, TITLE: Recovering Realistic Texture in Image Super-Resolution by Deep Spatial Feature Transform
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Recovering_Realistic_Texture_CVPR_2018_paper.html
AUTHORS: Xintao Wang, Ke Yu, Chao Dong, Chen Change Loy
HIGHLIGHT: In this paper, we show that it is possible to recover textures faithful to semantic classes.
- 64, TITLE: Deep Extreme Cut: From Extreme Points to Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Maninis_Deep_Extreme_Cut_CVPR_2018_paper.html
AUTHORS: Kevis-Kokitsi Maninis, Sergi Caelles, Jordi Pont-Tuset, Luc Van Gool
HIGHLIGHT: This paper explores the use of extreme points in an object (left-most, right-most, top, bottom pixels) as input to obtain precise object segmentation for images and videos.
- 65, TITLE: Learning to Parse Wireframes in Images of Man-Made Environments
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Learning_to_Parse_CVPR_2018_paper.html
AUTHORS: Kun Huang, Yifan Wang, Zihan Zhou, Tianjiao Ding, Shenghua Gao, Yi Ma
HIGHLIGHT: In this paper, we propose a learning-based approach to the task of automatically extracting a "wireframe" representation for images of cluttered man-made environments.
To this end, we have built a very large new dataset of over 5,000 images with wireframes thoroughly labelled by humans.
- 66, TITLE: Occlusion-Aware Rolling Shutter Rectification of 3D Scenes
http://openaccess.thecvf.com/content_cvpr_2018/html/Vasu_Occlusion-Aware_Rolling_Shutter_CVPR_2018_paper.html
AUTHORS: Subeesh Vasu, Mahesh Mohan M. R., A. N. Rajagopalan
HIGHLIGHT: In this work, we consider the specific scenario of a fast moving camera wherein the rolling shutter distortions not only are predominant but also become depth-dependent which in turn results in intra-frame occlusions.
- 67, TITLE: Content-Sensitive Supervoxels via Uniform Tessellations on Video Manifolds
http://openaccess.thecvf.com/content_cvpr_2018/html/Yi_Content-Sensitive_Supervoxels_via_CVPR_2018_paper.html
AUTHORS: Ran Yi, Yong-Jin Liu, Yu-Kun Lai
HIGHLIGHT: In this paper, we propose content-sensitive supervoxels (CSS), which are regularly-shaped 3D primitive volumes that possess the following characteristic: they are typically larger and longer in content-sparse regions (i.e., with homogeneous appearance and motion), and smaller and shorter in content-dense regions (i.e., with high variation of appearance and/or motion).
- 68, TITLE: Intrinsic Image Transformation via Scale Space Decomposition
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_Intrinsic_Image_Transformation_CVPR_2018_paper.html
AUTHORS: Lechao Cheng, Chengyi Zhang, Zicheng Liao
HIGHLIGHT: We introduce a new network structure for decomposing an image into its intrinsic albedo and shading.
- 69, TITLE: Learned Shape-Tailored Descriptors for Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Khan_Learned_Shape-Tailored_Descriptors_CVPR_2018_paper.html
AUTHORS: Naeemullah Khan, Ganesh Sundaramoorthi
HIGHLIGHT: We introduce and construct learned Shape-Tailored Descriptors that aggregate image statistics only within regions of interest to avoid mixing statistics of different textures, and that are invariant to complex nuisances (e.g., illumination, perspective and deformations).

70, TITLE: PAD-Net: Multi-Tasks Guided Prediction-and-Distillation Network for Simultaneous Depth Estimation and Scene Parsing
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_PAD-Net_Multi-Tasks_Guided_CVPR_2018_paper.html
AUTHORS: Dan Xu, Wanli Ouyang, Xiaogang Wang, Nicu Sebe
HIGHLIGHT: In this paper we tackle the problem of simultaneous depth estimation and scene parsing in a joint CNN.

71, TITLE: Multi-Image Semantic Matching by Mining Consistent Features
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Multi-Image_Semantic_Matching_CVPR_2018_paper.html
AUTHORS: Qianqian Wang, Xiaowei Zhou, Kostas Daniilidis
HIGHLIGHT: This work proposes a multi-image matching method to estimate semantic correspondences across multiple images.

72, TITLE: Density-Aware Single Image De-Raining Using a Multi-Stream Dense Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Density-Aware_Single_Image_CVPR_2018_paper.html
AUTHORS: He Zhang, Vishal M. Patel
HIGHLIGHT: We present a novel density-aware multi-stream densely connected convolutional neural network-based algorithm, called DID-MDN, for joint rain density estimation and de-raining.

73, TITLE: Joint Cuts and Matching of Partitions in One Graph
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Joint_Cuts_and_CVPR_2018_paper.html
AUTHORS: Tianshu Yu, Junchi Yan, Jieyi Zhao, Baoxin Li
HIGHLIGHT: Joint Cuts and Matching of Partitions in One Graph

74, TITLE: Progressive Attention Guided Recurrent Network for Salient Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Progressive_Attention_Guided_CVPR_2018_paper.html
AUTHORS: Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, Gang Wang
HIGHLIGHT: In this paper, we propose a novel attention guided network which selectively integrates multi-level contextual information in a progressive manner.

75, TITLE: Fast and Accurate Single Image Super-Resolution via Information Distillation Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Hui_Fast_and_Accurate_CVPR_2018_paper.html
AUTHORS: Zheng Hui, Xiumei Wang, Xinbo Gao
HIGHLIGHT: In order to solve the above questions, we propose a deep but compact convolutional network to directly reconstruct the high resolution image from the original low resolution image.

76, TITLE: Hallucinated-IQA: No-Reference Image Quality Assessment via Adversarial Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_Hallucinated-IQA_No-Reference_Image_CVPR_2018_paper.html
AUTHORS: Kwan-Yee Lin, Guanxiang Wang
HIGHLIGHT: In this work, we propose a hallucination-guided quality regression network to address the issue.

77, TITLE: NAG: Network for Adversary Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Mopuri_NAG_Network_for_CVPR_2018_paper.html
AUTHORS: Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, R. Venkatesh Babu
HIGHLIGHT: In this paper, we propose for the first time, a generative approach to model the distribution of adversarial perturbations.

78, TITLE: Dynamic-Structured Semantic Propagation Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Liang_Dynamic-Structured_Semantic_Propagation_CVPR_2018_paper.html
AUTHORS: Xiaodan Liang, Hongfei Zhou, Eric Xing
HIGHLIGHT: In this paper, we propose a Dynamic-Structured Semantic Propagation Network (DSSPN) that builds a semantic neuron graph to explicitly incorporate the concept hierarchy into dynamic network construction, leading to an interpretable reasoning process.

79, TITLE: Cross-Domain Self-Supervised Multi-Task Feature Learning Using Synthetic Imagery
http://openaccess.thecvf.com/content_cvpr_2018/html/Ren_Cross-Domain_Self-Supervised_Multi-Task_CVPR_2018_paper.html
AUTHORS: Zhongzheng Ren, Yong Jae Lee
HIGHLIGHT: In this paper, we propose a novel multi-task deep network to learn generalizable high-level visual representations.

80, TITLE: A Two-Step Disentanglement Method

- http://openaccess.thecvf.com/content_cvpr_2018/html/Hadad_A_Two-Step_Disentanglement_CVPR_2018_paper.html
AUTHORS: Naama Hadad, Lior Wolf, Moni Shahar
HIGHLIGHT: We address the problem of disentanglement of factors that generate a given data into those that are correlated with the labeling and those that are not.
- 81, TITLE: Robust Facial Landmark Detection via a Fully-Convolutional Local-Global Context Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Merget_Robust_Facial_Landmark_CVPR_2018_paper.html
AUTHORS: Daniel Merget, Matthias Rock, Gerhard Rigoll
HIGHLIGHT: In this work, we propose a new approach that introduces global context into a fully-convolutional neural network directly.
- 82, TITLE: Decorrelated Batch Normalization
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Decorrelated_Batch_Normalization_CVPR_2018_paper.html
AUTHORS: Lei Huang, Dawei Yang, Bo Lang, Jia Deng
HIGHLIGHT: In this work, we propose Decorrelated Batch Normalization (DBN), which not just centers and scales activations but whitens them.
- 83, TITLE: Learning to Sketch With Shortcut Cycle Consistency
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_Learning_to_Sketch_CVPR_2018_paper.html
AUTHORS: Jifei Song, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales
HIGHLIGHT: In this paper, we present a novel approach for translating an object photo to a sketch, mimicking the human sketching process.
- 84, TITLE: Towards a Mathematical Understanding of the Difficulty in Learning With Feedforward Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Towards_a_Mathematical_CVPR_2018_paper.html
AUTHORS: Hao Shen
HIGHLIGHT: This work provides an alternative mathematical understanding of the challenge from a smooth optimisation perspective.
- 85, TITLE: FaceID-GAN: Learning a Symmetry Three-Player GAN for Identity-Preserving Face Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_FaceID-GAN_Learning_a_CVPR_2018_paper.html
AUTHORS: Yujun Shen, Ping Luo, Junjie Yan, Xiaogang Wang, Xiaoou Tang
HIGHLIGHT: Existing methods typically formulate GAN as a two-player game, where a discriminator distinguishes face images from the real and synthesized domains, while a generator reduces its discriminativeness by synthesizing a face of photo-realistic quality.
- 86, TITLE: A Constrained Deep Neural Network for Ordinal Regression
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_A_Constrained_Deep_CVPR_2018_paper.html
AUTHORS: Yanzhu Liu, Adams Wai Kin Kong, Chi Keong Goh
HIGHLIGHT: This paper proposes a constrained optimization formulation for the ordinal regression problem which minimizes the negative loglikelihood for multiple categories constrained by the order relationship between instances.
- 87, TITLE: Modulated Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Modulated_Convolutional_Networks_CVPR_2018_paper.html
AUTHORS: Xiaodi Wang, Baochang Zhang, Ce Li, Rongrong Ji, Jungong Han, Xianbin Cao, Jianzhuang Liu
HIGHLIGHT: In this paper, we propose new Modulated Convolutional Networks (MCNs) to improve the portability of CNNs via binarized filters.
- 88, TITLE: Learning Steerable Filters for Rotation Equivariant CNNs
http://openaccess.thecvf.com/content_cvpr_2018/html/Weiler_Learning_Steerable_Filters_CVPR_2018_paper.html
AUTHORS: Maurice Weiler, Fred A. Hamprecht, Martin Storath
HIGHLIGHT: In this work, we develop Steerable Filter CNNs (SFCNNs) which achieve joint equivariance under translations and rotations by design.
- 89, TITLE: Efficient Interactive Annotation of Segmentation Datasets With Polygon-RNN++
http://openaccess.thecvf.com/content_cvpr_2018/html/Acuna_Efficient_Interactive_Annotation_CVPR_2018_paper.html
AUTHORS: David Acuna, Huan Ling, Amlan Kar, Sanja Fidler
HIGHLIGHT: In this work, we follow the idea of Polygon-RNN to produce polygonal annotations of objects interactively using humans-in-the-loop.

- 90, TITLE: SplineCNN: Fast Geometric Deep Learning With Continuous B-Spline Kernels
http://openaccess.thecvf.com/content_cvpr_2018/html/Fey_SplineCNN_Fast_Geometric_CVPR_2018_paper.html
AUTHORS: Matthias Fey, Jan Eric Lenssen, Frank Weichert, Heinrich Müller
HIGHLIGHT: We present Spline-based Convolutional Neural Networks (SplineCNNs), a variant of deep neural networks for irregular structured and geometric input, e.g., graphs or meshes.
- 91, TITLE: GAGAN: Geometry-Aware Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Kossaiifi_GAGAN_Geometry-Aware_Generative_CVPR_2018_paper.html
AUTHORS: Jean Kossaiifi, Linh Tran, Yannis Panagakis, Maja Pantic
HIGHLIGHT: This paper introduces the Geometry-Aware Generative Adversarial Networks (GAGAN) for incorporating geometric information into the image generation process.
- 92, TITLE: On the Robustness of Semantic Segmentation Models to Adversarial Attacks
http://openaccess.thecvf.com/content_cvpr_2018/html/Arnab_On_the_Robustness_CVPR_2018_paper.html
AUTHORS: Anurag Arnab, Ondrej Miksik, Philip H.S. Torr
HIGHLIGHT: In this paper, we present what to our knowledge is the first rigorous evaluation of adversarial attacks on modern semantic segmentation models, using two large-scale datasets.
- 93, TITLE: Feedback-Prop: Convolutional Neural Network Inference Under Partial Evidence
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Feedback-Prop_Convolutional_Neural_CVPR_2018_paper.html
AUTHORS: Tianlu Wang, Kota Yamaguchi, Vicente Ordonez
HIGHLIGHT: We propose an inference procedure for deep convolutional neural networks (CNNs) when partial evidence is available.
- 94, TITLE: Super-Resolving Very Low-Resolution Face Images With Supplementary Attributes
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Super-Resolving_Very_Low-Resolution_CVPR_2018_paper.html
AUTHORS: Xin Yu, Basura Fernando, Richard Hartley, Fatih Porikli
HIGHLIGHT: Given a tiny face image, conventional face hallucination methods aim to super-resolve its high-resolution (HR) counterpart by learning a mapping from an exemplar dataset.
- 95, TITLE: Frustum PointNets for 3D Object Detection From RGB-D Data
http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Frustum_PointNets_for_CVPR_2018_paper.html
AUTHORS: Charles R. Qi, Wei Liu, Chenxia Wu, Hao Su, Leonidas J. Guibas
HIGHLIGHT: In this work, we study 3D object detection from RGB-D data in both indoor and outdoor scenes.
- 96, TITLE: W2F: A Weakly-Supervised to Fully-Supervised Framework for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_W2F_A_Weakly-Supervised_CVPR_2018_paper.html
AUTHORS: Yongqiang Zhang, Yancheng Bai, Mingli Ding, Yongqiang Li, Bernard Ghanem
HIGHLIGHT: To overcome these issues, we propose a weakly-supervised to fully-supervised framework, where a weakly-supervised detector is implemented using multiple instance learning.
- 97, TITLE: 3D Object Detection With Latent Support Surfaces
http://openaccess.thecvf.com/content_cvpr_2018/html/Ren_3D_Object_Detection_CVPR_2018_paper.html
AUTHORS: Zhile Ren, Erik B. Sudderth
HIGHLIGHT: We develop a 3D object detection algorithm that uses latent support surfaces to capture contextual relationships in indoor scenes.
- 98, TITLE: Towards Faster Training of Global Covariance Pooling Networks by Iterative Matrix Square Root Normalization
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Towards_Faster_Training_CVPR_2018_paper.html
AUTHORS: Peihua Li, Jiangtao Xie, Qilong Wang, Zilin Gao
HIGHLIGHT: Towards addressing this problem, we propose an iterative matrix square root normalization method for fast end-to-end training of global covariance pooling networks.
- 99, TITLE: Recurrent Scene Parsing With Perspective Understanding in the Loop
http://openaccess.thecvf.com/content_cvpr_2018/html/Kong_Recurrent_Scene_Parsing_CVPR_2018_paper.html
AUTHORS: Shu Kong, Charless C. Fowlkes
HIGHLIGHT: We propose a depth-aware gating module that adaptively selects the pooling field size in a convolutional network architecture according to the object scale (inversely proportional to the depth) so that small details are preserved for distant objects while larger receptive fields are used for those nearby.

100, TITLE: Improving Occlusion and Hard Negative Handling for Single-Stage Pedestrian Detectors
http://openaccess.thecvf.com/content_cvpr_2018/html/Noh_Improving_Occlusion_and_CVPR_2018_paper.html
AUTHORS: Junhyug Noh, Soochan Lee, Beomsu Kim, Gunhee Kim
HIGHLIGHT: We propose methods of addressing two critical issues of pedestrian detection: (i) occlusion of target objects as false negative failure, and (ii) confusion with hard negative examples like vertical structures as false positive failure.

101, TITLE: Learning to Act Properly: Predicting and Explaining Affordances From Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Chuang_Learning_to_Act_CVPR_2018_paper.html
AUTHORS: Ching-Yao Chuang, Jiaman Li, Antonio Torralba, Sanja Fidler
HIGHLIGHT: We propose a model that exploits Graph Neural Networks to propagate contextual information from the scene in order to perform detailed affordance reasoning about each object.
We collect a new dataset that builds upon ADE20k, referred to as ADE-Affordance, which containing annotations enabling such rich visual reasoning.

102, TITLE: Pointwise Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Hua_Pointwise_Convolutional_Neural_CVPR_2018_paper.html
AUTHORS: Binh-Son Hua, Minh-Khoi Tran, Sai-Kit Yeung
HIGHLIGHT: In this paper, we present a convolutional neural network for semantic segmentation and object recognition with 3D point clouds.

103, TITLE: Image-Image Domain Adaptation With Preserved Self-Similarity and Domain-Dissimilarity for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Deng_Image-Image_Domain_Adaptation_CVPR_2018_paper.html
AUTHORS: Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, Jianbin Jiao
HIGHLIGHT: In our attempt, we present a "learning via translation" framework.

104, TITLE: A Generative Adversarial Approach for Zero-Shot Learning From Noisy Texts
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_A_Generative_Adversarial_CVPR_2018_paper.html
AUTHORS: Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, Ahmed Elgammal
HIGHLIGHT: Specifically, we propose a simple yet effective generative model that takes as input noisy text descriptions about an unseen class (e.g. Wikipedia articles) and generates synthesized visual features for this class.

105, TITLE: Tensorize, Factorize and Regularize: Robust Visual Relationship Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Hwang_Tensorize_Factorize_and_CVPR_2018_paper.html
AUTHORS: Seong Jae Hwang, Sathya N. Ravi, Zirui Tao, Hyunwoo J. Kim, Maxwell D. Collins, Vikas Singh
HIGHLIGHT: In this work, we start from a simple multi-relational learning model, which in principle, offers a rich formalization for deriving a strong prior for learning visual relationships.

106, TITLE: Transductive Unbiased Embedding for Zero-Shot Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_Transductive_Unbiased_Embedding_CVPR_2018_paper.html
AUTHORS: Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, Mingli Song
HIGHLIGHT: In this paper, we propose a straightforward yet effective method named Quasi-Fully Supervised Learning (QFSL) to alleviate the bias problem.

107, TITLE: Hierarchical Novelty Detection for Visual Object Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_Hierarchical_Novelty_Detection_CVPR_2018_paper.html
AUTHORS: Kibok Lee, Kimin Lee, Kyle Min, Yuting Zhang, Jinwoo Shin, Honglak Lee
HIGHLIGHT: In this paper, we study more informative novelty detection schemes based on a hierarchical classification framework.

108, TITLE: Zero-Shot Visual Recognition Using Semantics-Preserving Adversarial Embedding Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Zero-Shot_Visual_Recognition_CVPR_2018_paper.html
AUTHORS: Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, Shih-Fu Chang
HIGHLIGHT: We propose a novel framework called Semantics-Preserving Adversarial Embedding Network (SP-AEN) for zero-shot visual recognition (ZSL), where test images and their classes are both unseen during training.

109, TITLE: Learning Rich Features for Image Manipulation Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Learning_Rich_Features_CVPR_2018_paper.html
AUTHORS: Peng Zhou, Xintong Han, Vlad I. Morariu, Larry S. Davis

HIGHLIGHT: We propose a two-stream Faster R-CNN network and train it end-to-end to detect the tampered regions given a manipulated image.

110, **TITLE:** Human Semantic Parsing for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Kalayeh_Human_Semantic_Parsing_CVPR_2018_paper.html
AUTHORS: Mahdi M. Kalayeh, Emrah Basaran, Muhittin Gökmen, Mustafa E. Kamasak, Mubarak Shah
HIGHLIGHT: In this paper, we propose to adopt human semantic parsing which, due to its pixel-level accuracy and capability of modeling arbitrary contours, is naturally a better alternative.

111, **TITLE:** Stacked Latent Attention for Multimodal Reasoning
http://openaccess.thecvf.com/content_cvpr_2018/html/Fan_Stacked_Latent_Attention_CVPR_2018_paper.html
AUTHORS: Haoqi Fan, Jiatong Zhou
HIGHLIGHT: In this work, we pinpoint the potential limitations to the design of a traditional attention model.

112, **TITLE:** R-FCN-3000 at 30fps: Decoupling Detection and Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Singh_R-FCN-3000_at_30fps_CVPR_2018_paper.html
AUTHORS: Bharat Singh, Hengduo Li, Abhishek Sharma, Larry S. Davis
HIGHLIGHT: We propose a modular approach towards large-scale real-time object detection by decoupling objectness detection and classification.

113, **TITLE:** CSRNet: Dilated Convolutional Neural Networks for Understanding the Highly Congested Scenes
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_CSRNet_Dilated_Convolutional_CVPR_2018_paper.html
AUTHORS: Yuhong Li, Xiaofan Zhang, Deming Chen
HIGHLIGHT: We propose a network for Congested Scene Recognition called CSRNet to provide a data-driven and deep learning method that can understand highly congested scenes and perform accurate count estimation as well as present high-quality density maps.

114, **TITLE:** Revisiting Knowledge Transfer for Training Object Class Detectors
http://openaccess.thecvf.com/content_cvpr_2018/html/Uijlings_Revisiting_Knowledge_Transfer_CVPR_2018_paper.html
AUTHORS: Jasper Uijlings, Stefan Popov, Vittorio Ferrari
HIGHLIGHT: We present a unified knowledge transfer framework based on training a single neural network multi-class object detector over all source classes, organized in a semantic hierarchy.

115, **TITLE:** Deep Sparse Coding for Invariant Multimodal Halle Berry Neurons
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_Deep_Sparse_Coding_CVPR_2018_paper.html
AUTHORS: Edward Kim, Darryl Hannan, Garrett Kenyon
HIGHLIGHT: We define our model as a sparse coding problem using hierarchical layers.

116, **TITLE:** On the Convergence of PatchMatch and Its Variants
http://openaccess.thecvf.com/content_cvpr_2018/html/Ehret_On_the_Convergence_CVPR_2018_paper.html
AUTHORS: Thibaud Ehret, Pablo Arias
HIGHLIGHT: In this work we propose a theoretical framework for the analysis of PatchMatch and its variants, and apply it to derive bounds on their convergence rate.

117, **TITLE:** Rethinking the Faster R-CNN Architecture for Temporal Action Localization
http://openaccess.thecvf.com/content_cvpr_2018/html/Chao_Rethinking_the_Faster_CVPR_2018_paper.html
AUTHORS: Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, Rahul Sukthankar
HIGHLIGHT: We propose TAL-Net, an improved approach to temporal action localization in video that is inspired by the Faster R-CNN object detection framework.

118, **TITLE:** MoNet: Deep Motion Exploitation for Video Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Xiao_MoNet_Deep_Motion_CVPR_2018_paper.html
AUTHORS: Huaxin Xiao, Jiashi Feng, Guosheng Lin, Yu Liu, Maojun Zhang
HIGHLIGHT: In this paper, we propose a novel MoNet model to deeply exploit motion cues for boosting video object segmentation performance from two aspects, i.e., frame representation learning and segmentation refinement.

119, **TITLE:** Video Representation Learning Using Discriminative Pooling
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Video_Representation_Learning_CVPR_2018_paper.html
AUTHORS: Jue Wang, Anoop Cherian, Fatih Porikli, Stephen Gould

HIGHLIGHT: In an attempt to tackle this problem, we propose discriminative pooling, based on the notion that among the deep features generated on all short clips, there is at least one that characterizes the action.

120, **TITLE:** Recognizing Human Actions as the Evolution of Pose Estimation Maps
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Recognizing_Human_Actions_CVPR_2018_paper.html
AUTHORS: Mengyuan Liu, Junsong Yuan
HIGHLIGHT: With recent advances of human pose estimation, this work presents a novel method to recognize human action as the evolution of pose estimation maps.

121, **TITLE:** Video Person Re-Identification With Competitive Snippet-Similarity Aggregation and Co-Attentive Snippet Embedding
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Video_Person_Re-Identification_CVPR_2018_paper.html
AUTHORS: Dapeng Chen, Hongsheng Li, Tong Xiao, Shuai Yi, Xiaogang Wang
HIGHLIGHT: In this paper, we address video-based person re-identification with competitive snippet-similarity aggregation and co-attentive snippet embedding.

122, **TITLE:** Mask-Guided Contrastive Attention Model for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_Mask-Guided_Contrastive_Attention_CVPR_2018_paper.html
AUTHORS: Chunfeng Song, Yan Huang, Wanli Ouyang, Liang Wang
HIGHLIGHT: We evaluate the proposed method on three public datasets, including MARS, Market-1501 and CUHK03.

123, **TITLE:** Blazingly Fast Video Object Segmentation With Pixel-Wise Metric Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Blazingly_Fast_Video_CVPR_2018_paper.html
AUTHORS: Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, Luc Van Gool
HIGHLIGHT: This paper tackles the problem of video object segmentation, given some user annotation which indicates the object of interest.

124, **TITLE:** Learning to Compare: Relation Network for Few-Shot Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Sung_Learning_to_Compare_CVPR_2018_paper.html
AUTHORS: Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, Timothy M. Hospedales
HIGHLIGHT: We present a conceptually simple, flexible, and general framework for few-shot learning, where a classifier must learn to recognise new classes given only few examples from each.

125, **TITLE:** COCO-Stuff: Thing and Stuff Classes in Context
http://openaccess.thecvf.com/content_cvpr_2018/html/Caesar_COCO-Stuff_Thing_and_CVPR_2018_paper.html
AUTHORS: Holger Caesar, Jasper Uijlings, Vittorio Ferrari
HIGHLIGHT: To understand stuff and things in context we introduce COCO-Stuff, which augments all 164K images of the COCO 2017 dataset with pixel-wise annotations for 91 stuff classes.

126, **TITLE:** Image Generation From Scene Graphs
http://openaccess.thecvf.com/content_cvpr_2018/html/Johnson_Image_Generation_From_CVPR_2018_paper.html
AUTHORS: Justin Johnson, Agrim Gupta, Li Fei-Fei
HIGHLIGHT: To overcome this limitation we propose a method for generating images from scene graphs, enabling explicitly reasoning about objects and their relationships.

127, **TITLE:** Deep Cauchy Hashing for Hamming Space Retrieval
http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Deep_Cauchy_Hashing_CVPR_2018_paper.html
AUTHORS: Yue Cao, Mingsheng Long, Bin Liu, Jianmin Wang
HIGHLIGHT: This work presents Deep Cauchy Hashing (DCH), a novel deep hashing model that generates compact and concentrated binary hash codes to enable efficient and effective Hamming space retrieval.

128, **TITLE:** Learning to Look Around: Intelligently Exploring Unseen Environments for Unknown Tasks
http://openaccess.thecvf.com/content_cvpr_2018/html/Jayaraman_Learning_to_Look_CVPR_2018_paper.html
AUTHORS: Dinesh Jayaraman, Kristen Grauman
HIGHLIGHT: We propose a reinforcement learning solution, where the agent is rewarded for actions that reduce its uncertainty about the unobserved portions of its environment.

129, **TITLE:** Multi-Scale Location-Aware Kernel Representation for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Multi-Scale_Location-Aware_Kernel_CVPR_2018_paper.html
AUTHORS: Hao Wang, Qilong Wang, Mingqi Gao, Peihua Li, Wangmeng Zuo

HIGHLIGHT: In this paper, we make an attempt to exploit high-order statistics in object detection, aiming at generating more discriminative representations for proposals to enhance the performance of detectors.

130, **TITLE:** Clinical Skin Lesion Diagnosis Using Representations Inspired by Dermatologist Criteria
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Clinical_Skin_Lesion_CVPR_2018_paper.html

AUTHORS: Jufeng Yang, Xiaoxiao Sun, Jie Liang, Paul L. Rosin

HIGHLIGHT: To address this problem, we present effective representations inspired by the accepted dermatological criteria for diagnosing clinical skin lesions.

131, **TITLE:** Compare and Contrast: Learning Prominent Visual Differences

http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Compare_and_Contrast_CVPR_2018_paper.html

AUTHORS: Steven Chen, Kristen Grauman

HIGHLIGHT: In this work, we introduce and model prominent differences, a rich new functionality for comparing images.

132, **TITLE:** Multi-Evidence Filtering and Fusion for Multi-Label Classification, Object Detection and Semantic Segmentation Based on Weakly Supervised Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Ge_Multi-Evidence_Filtering_and_CVPR_2018_paper.html

AUTHORS: Weifeng Ge, Sibe Yang, Yizhou Yu

HIGHLIGHT: In this paper, we propose a novel weakly supervised curriculum learning pipeline for multi-label object recognition, detection and semantic segmentation.

133, **TITLE:** HashGAN: Deep Learning to Hash With Pair Conditional Wasserstein GAN

http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_HashGAN_Deep_Learning_CVPR_2018_paper.html

AUTHORS: Yue Cao, Bin Liu, Mingsheng Long, Jianmin Wang

HIGHLIGHT: This paper presents HashGAN, a novel architecture for deep learning to hash, which learns compact binary hash codes from both real images and diverse images synthesized by generative models.

134, **TITLE:** Min-Entropy Latent Model for Weakly Supervised Object Detection

http://openaccess.thecvf.com/content_cvpr_2018/html/Wan_Min-Entropy_Latent_Model_CVPR_2018_paper.html

AUTHORS: Fang Wan, Pengxu Wei, Jianbin Jiao, Zhenjun Han, Qixiang Ye

HIGHLIGHT: In this paper, a min-entropy latent model (MELM) is proposed for weakly supervised object detection.

135, **TITLE:** MAttNet: Modular Attention Network for Referring Expression Comprehension

http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_MAttNet_Modular_Attention_CVPR_2018_paper.html

AUTHORS: Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, Tamara L. Berg

HIGHLIGHT: In this paper, we address referring expression comprehension: localizing an image region described by a natural language expression.

136, **TITLE:** AttnGAN: Fine-Grained Text to Image Generation With Attentional Generative Adversarial Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_AttnGAN_Fine-Grained_Text_CVPR_2018_paper.html

AUTHORS: Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, Xiaodong He

HIGHLIGHT: In this paper, we propose an Attentional Generative Adversarial Network (AttnGAN) that allows attention-driven, multi-stage refinement for fine-grained text-to-image generation.

137, **TITLE:** Adversarial Complementary Learning for Weakly Supervised Object Localization

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Adversarial_Complementary_Learning_CVPR_2018_paper.html

AUTHORS: Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, Thomas S. Huang

HIGHLIGHT: In this work, we propose Adversarial Complementary Learning (ACoL) to automatically localize integral objects of semantic interest with weak supervision.

138, **TITLE:** Conditional Generative Adversarial Network for Structured Domain Adaptation

http://openaccess.thecvf.com/content_cvpr_2018/html/Hong_Conditional_Generative_Adversarial_CVPR_2018_paper.html

AUTHORS: Weixiang Hong, Zhenzhen Wang, Ming Yang, Junsong Yuan

HIGHLIGHT: In this paper, we propose a principled way to conduct structured domain adaption for semantic segmentation, i.e., integrating GAN into the FCN framework to mitigate the gap between source and target domains.

139, **TITLE:** GroupCap: Group-Based Image Captioning With Structured Relevance and Diversity Constraints

http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_GroupCap_Group-Based_Image_CVPR_2018_paper.html

AUTHORS: Fuhai Chen, Rongrong Ji, Xiaoshuai Sun, Yongjian Wu, Jinsong Su

HIGHLIGHT: In this paper, we propose a novel group-based image captioning scheme (termed GroupCap), which jointly models the structured relevance and diversity among visual contents of group images towards an optimal collaborative captioning. To facilitate quantitative evaluation, we further release two group captioning datasets derived from the MS-COCO benchmark, serving as the first of their kind.

140, **TITLE:** Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Weakly-Supervised_Semantic_Segmentation_CVPR_2018_paper.html
AUTHORS: Xiang Wang, Shaodi You, Xi Li, Huimin Ma
HIGHLIGHT: To bridge this gap, in this paper, we propose an iterative bottom-up and top-down framework which alternatively expands object regions and optimizes segmentation network.

141, **TITLE:** Bootstrapping the Performance of Weakly Supervised Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Bootstrapping_the_Performance_CVPR_2018_paper.html
AUTHORS: Tong Shen, Guosheng Lin, Chunhua Shen, Ian Reid
HIGHLIGHT: In this work, we focus on weak supervision, developing a method for training a high-quality pixel-level classifier for semantic segmentation, using only image-level class labels as the provided ground-truth.

142, **TITLE:** DeepVoting: A Robust and Explainable Deep Network for Semantic Part Detection Under Partial Occlusion
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_DeepVoting_A_Robust_CVPR_2018_paper.html
AUTHORS: Zhishuai Zhang, Cihang Xie, Jianyu Wang, Lingxi Xie, Alan L. Yuille
HIGHLIGHT: In this paper, we study the task of detecting semantic parts of an object, e.g., a wheel of a car, under partial occlusion.

143, **TITLE:** Geometry-Aware Scene Text Detection With Instance Transformation Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Geometry-Aware_Scene_Text_CVPR_2018_paper.html
AUTHORS: Fangfang Wang, Liming Zhao, Xi Li, Xinchao Wang, Dacheng Tao
HIGHLIGHT: In this paper, we propose a geometry-aware modeling approach tailored for scene text representation with an end-to-end learning scheme.

144, **TITLE:** Optical Flow Guided Feature: A Fast and Robust Motion Representation for Video Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Optical_Flow_Guided_CVPR_2018_paper.html
AUTHORS: Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, Wei Zhang
HIGHLIGHT: In this study, we introduce a novel compact motion representation for video action recognition, named Optical Flow guided Feature (OFF), which enables the network to distill temporal information through a fast and robust approach.

145, **TITLE:** Motion-Guided Cascaded Refinement Network for Video Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Motion-Guided_Cascaded_Refinement_CVPR_2018_paper.html
AUTHORS: Ping Hu, Gang Wang, Xiangfei Kong, Jason Kuen, Yap-Peng Tan
HIGHLIGHT: To tackle this problem, we propose a Motion-guided Cascaded Refinement Network for VOS.

146, **TITLE:** A Memory Network Approach for Story-Based Temporal Summarization of 360° Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_A_Memory_Network_CVPR_2018_paper.html
AUTHORS: Sangho Lee, Jinyoung Sung, Youngjae Yu, Gunhee Kim
HIGHLIGHT: We propose a novel memory network model named Past-Future Memory Network (PFMN), in which we first compute the scores of 81 normal field of view (NFOV) region proposals cropped from the input 360° video, and then recover a latent, collective summary using the network with two external memories that store the embeddings of previously selected subshots and future candidate subshots.

147, **TITLE:** Cube Padding for Weakly-Supervised Saliency Prediction in 360° Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_Cube_Padding_for_CVPR_2018_paper.html
AUTHORS: Hsien-Tzu Cheng, Chun-Hung Chao, Jin-Dong Dong, Hao-Kai Wen, Tyng-Luh Liu, Min Sun
HIGHLIGHT: In contrast, we propose a simple and effective Cube Padding (CP) technique as follows.

148, **TITLE:** Appearance-and-Relation Networks for Video Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Appearance-and-Relation_Networks_for_CVPR_2018_paper.html
AUTHORS: Limin Wang, Wei Li, Wen Li, Luc Van Gool
HIGHLIGHT: This paper presents a new architecture, termed as Appearance-and-Relation Network (ARTNet), to learn video representation in an end-to-end manner.

149, **TITLE:** Excitation Backprop for RNNs

http://openaccess.thecvf.com/content_cvpr_2018/html/Bargal_Excitation_Backprop_for_CVPR_2018_paper.html
AUTHORS: Sarah Adel Bargal, Andrea Zunino, Donghyun Kim, Jianming Zhang, Vittorio Murino, Stan Sclaroff
HIGHLIGHT: In this work, we devise a formulation that simultaneously grounds evidence in space and time, in a single pass, using top-down saliency.

150, TITLE: One-Shot Action Localization by Learning Sequence Matching Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_One-Shot_Action_Localization_CVPR_2018_paper.html
AUTHORS: Hongtao Yang, Xuming He, Fatih Porikli
HIGHLIGHT: Towards this objective, we introduce a novel one-shot action localization method that alleviates the need for large amounts of training samples.

151, TITLE: Structure Preserving Video Prediction
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Structure_Preserving_Video_CVPR_2018_paper.html
AUTHORS: Jingwei Xu, Bingbing Ni, Zefan Li, Shuo Cheng, Xiaokang Yang
HIGHLIGHT: To this end, we present a structure preserving video prediction framework to explicitly address above issues and enhance video prediction quality.

152, TITLE: Person Re-Identification With Cascaded Pairwise Convolutions
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Person_Re-Identification_With_CVPR_2018_paper.html
AUTHORS: Yicheng Wang, Zhenzhong Chen, Feng Wu, Gang Wang
HIGHLIGHT: In this paper, a novel deep architecture named BraidNet is proposed for person re-identification.

153, TITLE: On the Importance of Label Quality for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zlateski_On_the_Importance_CVPR_2018_paper.html
AUTHORS: Aleksandar Zlateski, Ronnachai Jaroensri, Prafull Sharma, Frédo Durand
HIGHLIGHT: We investigate the relationship between the quality of labels and the performance of ConvNets for semantic segmentation.
We create a very large synthetic dataset with perfectly labeled street view scenes.

154, TITLE: Scalable and Effective Deep CCA via Soft Decorrelation
http://openaccess.thecvf.com/content_cvpr_2018/html/Chang_Scalable_and_Effective_CVPR_2018_paper.html
AUTHORS: Xiaobin Chang, Tao Xiang, Timothy M. Hospedales
HIGHLIGHT: We propose a novel deep CCA model Soft CCA to overcome these problems.

155, TITLE: Duplex Generative Adversarial Network for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Duplex_Generative_Adversarial_CVPR_2018_paper.html
AUTHORS: Lanqing Hu, Meina Kan, Shiguang Shan, Xilin Chen
HIGHLIGHT: Following the similar idea of GAN, this work proposes a novel GAN architecture with duplex adversarial discriminators (referred to as DupGAN), which can achieve domain-invariant representation and domain transformation.

156, TITLE: Edit Probability for Scene Text Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Bai_Edit_Probability_for_CVPR_2018_paper.html
AUTHORS: Fan Bai, Zhanzhan Cheng, Yi Niu, Shiliang Pu, Shuigeng Zhou
HIGHLIGHT: To handle this problem, we propose a novel method called edit probability (EP) for scene text recognition.

157, TITLE: Global Versus Localized Generative Adversarial Nets
http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Global_Versus_Localized_CVPR_2018_paper.html
AUTHORS: Guo-Jun Qi, Liheng Zhang, Hao Hu, Marzieh Edraki, Jingdong Wang, Xian-Sheng Hua
HIGHLIGHT: In this paper, we present a novel localized Generative Adversarial Net (GAN) to learn on the manifold of real data.

158, TITLE: MoCoGAN: Decomposing Motion and Content for Video Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Tulyakov_MoCoGAN_Decomposing_Motion_CVPR_2018_paper.html
AUTHORS: Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, Jan Kautz
HIGHLIGHT: To learn motion and content decomposition in an unsupervised manner, we introduce a novel adversarial learning scheme utilizing both image and video discriminators.

159, TITLE: Recurrent Residual Module for Fast Inference in Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Pan_Recurrent_Residual_Module_CVPR_2018_paper.html
AUTHORS: Bowen Pan, Wuwei Lin, Xiaolin Fang, Chaoqin Huang, Bolei Zhou, Cewu Lu

HIGHLIGHT: In this work, we propose a framework called Recurrent Residual Module (RRM) to accelerate the CNN inference for video recognition tasks.

160, **TITLE:** Improving Landmark Localization With Semi-Supervised Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Honari_Improving_Landmark_Localization_CVPR_2018_paper.html
AUTHORS: Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, Jan Kautz
HIGHLIGHT: We present two techniques to improve landmark localization in images from partially annotated datasets.

161, **TITLE:** Adversarial Data Programming: Using GANs to Relax the Bottleneck of Curated Labeled Data
http://openaccess.thecvf.com/content_cvpr_2018/html/Pal_Adversarial_Data_Programming_CVPR_2018_paper.html
AUTHORS: Arghya Pal, Vineeth N. Balasubramanian
HIGHLIGHT: In this work, we present Adversarial Data Programming (ADP), which presents an adversarial methodology to generate data as well as a curated aggregated label, given a set of weak labeling functions.

162, **TITLE:** Stochastic Variational Inference With Gradient Linearization
http://openaccess.thecvf.com/content_cvpr_2018/html/Plotz_Stochastic_Variational_Inference_CVPR_2018_paper.html
AUTHORS: Tobias Plötz, Anne S. Wannenwetsch, Stefan Roth
HIGHLIGHT: In this paper we propose stochastic variational inference with gradient linearization (SVIGL).

163, **TITLE:** Multi-Label Zero-Shot Learning With Structured Knowledge Graphs
http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_Multi-Label_Zero-Shot_Learning_CVPR_2018_paper.html
AUTHORS: Chung-Wei Lee, Wei Fang, Chih-Kuan Yeh, Yu-Chiang Frank Wang
HIGHLIGHT: In this paper, we propose a novel deep learning architecture for multi-label zero-shot learning (ML-ZSL), which is able to predict multiple unseen class labels for each input instance.

164, **TITLE:** MorphNet: Fast & Simple Resource-Constrained Structure Learning of Deep Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Gordon_MorphNet_Fast_CVPR_2018_paper.html
AUTHORS: Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, Edward Choi
HIGHLIGHT: We present MorphNet, an approach to automate the design of neural network structures.

165, **TITLE:** Deep Adversarial Subspace Clustering
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Deep_Adversarial_Subspace_CVPR_2018_paper.html
AUTHORS: Pan Zhou, Yunqing Hou, Jiashi Feng
HIGHLIGHT: To solve this issue, we propose a novel deep adversarial subspace clustering (DASC) model, which learns more favorable sample representations by deep learning for subspace clustering, and more importantly introduces adversarial learning to supervise sample representation learning and subspace clustering.

166, **TITLE:** Towards Human-Machine Cooperation: Self-Supervised Sample Mining for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Towards_Human-Machine_Cooperation_CVPR_2018_paper.html
AUTHORS: Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, Liang Lin
HIGHLIGHT: In this paper, aiming to remedy the drawbacks of existing AL methods, we present a principled Self-supervised Sample Mining (SSM) process accounting for the real challenges in object detection.

167, **TITLE:** Discrete-Continuous ADMM for Transductive Inference in Higher-Order MRFs
http://openaccess.thecvf.com/content_cvpr_2018/html/Laude_Discrete-Continuous_ADMM_for_CVPR_2018_paper.html
AUTHORS: Emanuel Laude, Jan-Hendrik Lange, Jonas Schüpfer, Csaba Domokos, Laura Leal-Taixé, Frank R. Schmidt, Bjoern Andres, Daniel Cremers
HIGHLIGHT: This paper introduces a novel algorithm for transductive inference in higher-order MRFs, where the unary energies are parameterized by a variable classifier.

168, **TITLE:** Robust Physical-World Attacks on Deep Learning Visual Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Eykholt_Robust_Physical-World_Attacks_CVPR_2018_paper.html
AUTHORS: Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, Dawn Song
HIGHLIGHT: We propose a general attack algorithm, Robust Physical Perturbations (RP²), to generate robust visual adversarial perturbations under different physical conditions.

169, **TITLE:** Generating a Fusion Image: One's Identity and Another's Shape
http://openaccess.thecvf.com/content_cvpr_2018/html/Joo_Generating_a_Fusion_CVPR_2018_paper.html
AUTHORS: DongGyu Joo, Doyeon Kim, Junmo Kim

HIGHLIGHT: Generating a novel image by manipulating two input images is an interesting research problem in the study of generative adversarial networks (GANs).

170, **TITLE:** Learning to Promote Saliency Detectors
http://openaccess.thecvf.com/content_cvpr_2018/html/Zeng_Learning_to_Promote_CVPR_2018_paper.html
AUTHORS: Yu Zeng, Huchuan Lu, Lihe Zhang, Mengyang Feng, Ali Borji
HIGHLIGHT: To solve this issue, we formulate a zero-shot learning problem to promote existing saliency detectors.

171, **TITLE:** Image Super-Resolution via Dual-State Recurrent Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Han_Image_Super-Resolution_via_CVPR_2018_paper.html
AUTHORS: Wei Han, Shiyu Chang, Ding Liu, Mo Yu, Michael Witbrock, Thomas S. Huang
HIGHLIGHT: In this paper, we explore new structures for SR based on this compact RNN view, leading us to a dual-state design, the Dual-State Recurrent Network (DSRN).

172, **TITLE:** Deep Back-Projection Networks for Super-Resolution
http://openaccess.thecvf.com/content_cvpr_2018/html/Haris_Deep_Back-Projection_Networks_CVPR_2018_paper.html
AUTHORS: Muhammad Haris, Gregory Shakhnarovich, Norimichi Ukita
HIGHLIGHT: We propose Deep Back-Projection Networks (DBPN), that exploit iterative up- and down-sampling layers, providing an error feedback mechanism for projection errors at each stage.

173, **TITLE:** Focus Manipulation Detection via Photometric Histogram Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Focus_Manipulation_Detection_CVPR_2018_paper.html
AUTHORS: Can Chen, Scott McCloskey, Jingyi Yu
HIGHLIGHT: Classic image forensic methods leverage low-level cues such as metadata, sensor noise fingerprints, and others that are easily fooled when the image is re-encoded upon upload to facebook, etc.

174, **TITLE:** Compassionately Conservative Balanced Cuts for Image Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Cahill_Compassionately_Conservative_Balanced_CVPR_2018_paper.html
AUTHORS: Nathan D. Cahill, Tyler L. Hayes, Renee T. Meinhold, John F. Hamilton
HIGHLIGHT: We show that CCB-Cut minimization can be relaxed into an orthogonally constrained $\ell_{\{au\}}$ -minimization problem that coincides with the problem of computing Piecewise Flat Embeddings (PFE) for one particular index value, and we present an algorithm for solving the relaxed problem by iteratively minimizing a sequence of reweighted Rayleigh quotients (IRRQ).

175, **TITLE:** A High-Quality Denoising Dataset for Smartphone Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Abdelhamed_A_High-Quality_Denoising_CVPR_2018_paper.html
AUTHORS: Abdelrahman Abdelhamed, Stephen Lin, Michael S. Brown
HIGHLIGHT: We address this issue in this paper with the following contributions.

176, **TITLE:** Context-Aware Synthesis for Video Frame Interpolation
http://openaccess.thecvf.com/content_cvpr_2018/html/Niklaus_Context-Aware_Synthesis_for_CVPR_2018_paper.html
AUTHORS: Simon Niklaus, Feng Liu
HIGHLIGHT: This paper presents a context-aware synthesis approach that warps not only the input frames but also their pixel-wise contextual information and uses them to interpolate a high-quality intermediate frame.

177, **TITLE:** Salient Object Detection Driven by Fixation Prediction
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Salient_Object_Detection_CVPR_2018_paper.html
AUTHORS: Wenguan Wang, Jianbing Shen, Xingping Dong, Ali Borji
HIGHLIGHT: In this paper, we propose to employ the former model type to identify and segment salient objects in scenes.

178, **TITLE:** Enhancing the Spatial Resolution of Stereo Images Using a Parallax Prior
http://openaccess.thecvf.com/content_cvpr_2018/html/Jeon_Enhancing_the_Spatial_CVPR_2018_paper.html
AUTHORS: Daniel S. Jeon, Seung-Hwan Baek, Inchang Choi, Min H. Kim
HIGHLIGHT: We present a novel method that can enhance the spatial resolution of stereo images using a parallax prior.

179, **TITLE:** HATS: Histograms of Averaged Time Surfaces for Robust Event-Based Object Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Sironi_HATS_Histograms_of_CVPR_2018_paper.html
AUTHORS: Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, Ryad Benosman
HIGHLIGHT: In this paper we address both problems. Second, we release the first large real-world event-based dataset for object classification.

- 180, TITLE: A Bi-Directional Message Passing Model for Salient Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_A_Bi-Directional_Message_CVPR_2018_paper.html
AUTHORS: Lu Zhang, Ju Dai, Huchuan Lu, You He, Gang Wang
HIGHLIGHT: In this paper, we propose a novel bi-directional message passing model to integrate multi-level features for salient object detection.
- 181, TITLE: Matching Pixels Using Co-Occurrence Statistics
http://openaccess.thecvf.com/content_cvpr_2018/html/Kat_Matching_Pixels_Using_CVPR_2018_paper.html
AUTHORS: Rotal Kat, Roy Jevnisek, Shai Avidan
HIGHLIGHT: We propose a new error measure for matching pixels that is based on co-occurrence statistics.
- 182, TITLE: SeedNet: Automatic Seed Generation With Deep Reinforcement Learning for Robust Interactive Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_SeedNet_Automatic_Seed_CVPR_2018_paper.html
AUTHORS: Gwangmo Song, Heesoo Myeong, Kyoung Mu Lee
HIGHLIGHT: In this paper, we propose an automatic seed generation technique with deep reinforcement learning to solve the interactive segmentation problem.
- 183, TITLE: Jerk-Aware Video Acceleration Magnification
http://openaccess.thecvf.com/content_cvpr_2018/html/Takeda_Jerk-Aware_Video_Acceleration_CVPR_2018_paper.html
AUTHORS: Shoichiro Takeda, Kazuki Okami, Dan Mikami, Megumi Isogai, Hideaki Kimata
HIGHLIGHT: In this paper, we present a novel use of jerk to make the acceleration method robust to quick large motions.
- 184, TITLE: Defense Against Adversarial Attacks Using High-Level Representation Guided Denoiser
http://openaccess.thecvf.com/content_cvpr_2018/html/Liao_Defense_Against_Adversarial_CVPR_2018_paper.html
AUTHORS: Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, Jun Zhu
HIGHLIGHT: We propose high-level representation guided denoiser (HGD) as a defense for image classification.
- 185, TITLE: Stacked Conditional Generative Adversarial Networks for Jointly Learning Shadow Detection and Shadow Removal
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Stacked_Conditional_Generative_CVPR_2018_paper.html
AUTHORS: Jifeng Wang, Xiang Li, Jian Yang
HIGHLIGHT: In this paper, we present a multi-task perspective, which is not embraced by any existing work, to jointly learn both detection and removal in an end-to-end fashion that aims at enjoying the mutually improved benefits from each other. To fully evaluate the performance of our proposed framework, we construct the first large-scale benchmark with 1870 image triplets (shadow image, shadow mask image, and shadow-free image) under 135 scenes.
- 186, TITLE: Image Correction via Deep Reciprocating HDR Transformation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Image_Correction_via_CVPR_2018_paper.html
AUTHORS: Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, Rynson W.H. Lau
HIGHLIGHT: To this end, we propose an end-to-end DRHT model, which contains two CNNs, one for HDR detail reconstruction and the other for LDR detail correction.
- 187, TITLE: PieAPP: Perceptual Image-Error Assessment Through Pairwise Preference
http://openaccess.thecvf.com/content_cvpr_2018/html/Prashnani_PieAPP_Perceptual_Image-Error_CVPR_2018_paper.html
AUTHORS: Ekta Prashnani, Hong Cai, Yasamin Mostofi, Pradeep Sen
HIGHLIGHT: In this paper, we present a new learning-based method that is the first to predict perceptual image error like human observers. Since it is much easier for people to compare two given images and identify the one more similar to a reference than to assign quality scores to each, we propose a new, large-scale dataset labeled with the probability that humans will prefer one image over another.
- 188, TITLE: Normalized Cut Loss for Weakly-Supervised CNN Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Tang_Normalized_Cut_Loss_CVPR_2018_paper.html
AUTHORS: Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, Christopher Schroers
HIGHLIGHT: Inspired by the general ideas in semi-supervised learning, we address these problems via a new principled loss function evaluating network output with criteria standard in "shallow" segmentation, e.g. normalized cut.
- 189, TITLE: ISTA-Net: Interpretable Optimization-Inspired Deep Network for Image Compressive Sensing
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_ISTA-Net_Interpretable_Optimization-Inspired_CVPR_2018_paper.html
AUTHORS: Jian Zhang, Bernard Ghanem

HIGHLIGHT: Specifically, we propose a novel structured deep network, dubbed ISTA-Net, which is inspired by the Iterative Shrinkage-Thresholding Algorithm (ISTA) for optimizing a general L1 norm CS reconstruction model.

190, **TITLE:** Fast End-to-End Trainable Guided Filter
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Fast_End-to-End_Trainable_CVPR_2018_paper.html
AUTHORS: Huikai Wu, Shuai Zheng, Junge Zhang, Kaiqi Huang
HIGHLIGHT: We present a deep learning building block for joint upsampling, namely guided filtering layer.

191, **TITLE:** Disentangling Structure and Aesthetics for Style-Aware Image Completion
http://openaccess.thecvf.com/content_cvpr_2018/html/Gilbert_Disentangling_Structure_and_CVPR_2018_paper.html
AUTHORS: Andrew Gilbert, John Collomosse, Hailin Jin, Brian Price
HIGHLIGHT: We propose a non-parametric in-painting algorithm that enforces both structural and aesthetic (style) consistency within the resulting image.

192, **TITLE:** Learning a Discriminative Feature Network for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Learning_a_Discriminative_CVPR_2018_paper.html
AUTHORS: Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, Nong Sang
HIGHLIGHT: To tackle these two problems, we propose a Discriminative Feature Network (DFN), which contains two sub-networks: Smooth Network and Border Network.

193, **TITLE:** Kernelized Subspace Pooling for Deep Local Descriptors
http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Kernelized_Subspace_Pooling_CVPR_2018_paper.html
AUTHORS: Xing Wei, Yue Zhang, Yihong Gong, Nanning Zheng
HIGHLIGHT: In this paper, we propose a descriptor that has both highly invariant and discriminative power.

194, **TITLE:** pOSE: Pseudo Object Space Error for Initialization-Free Bundle Adjustment
http://openaccess.thecvf.com/content_cvpr_2018/html/Hong_pOSE_Pseudo_Object_CVPR_2018_paper.html
AUTHORS: Je Hyeong Hong, Christopher Zach
HIGHLIGHT: In this paper, we propose the Pseudo Object Space Error (pOSE), which is an objective with cameras represented as a hybrid between the affine and projective models.

195, **TITLE:** Deformable Shape Completion With Graph Convolutional Autoencoders
http://openaccess.thecvf.com/content_cvpr_2018/html/Litany_Deformable_Shape_Completion_CVPR_2018_paper.html
AUTHORS: Or Litany, Alex Bronstein, Michael Bronstein, Ameesh Makadia
HIGHLIGHT: In this work, we propose a novel learningbased method for the completion of partial shapes.

196, **TITLE:** Learning From Millions of 3D Scans for Large-Scale 3D Face Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Gilani_Learning_From_Millions_CVPR_2018_paper.html
AUTHORS: Syed Zulqarnain Gilani, Ajmal Mian
HIGHLIGHT: In this backdrop, we propose a method for generating a large corpus of labeled 3D face identities and their multiple instances for training and a protocol for merging the most challenging existing 3D datasets for testing.

197, **TITLE:** CarFusion: Combining Point Tracking and Part Detection for Dynamic 3D Reconstruction of Vehicles
http://openaccess.thecvf.com/content_cvpr_2018/html/Reddy_CarFusion_Combining_Point_CVPR_2018_paper.html
AUTHORS: N. Dinesh Reddy, Minh Vo, Srinivasa G. Narasimhan
HIGHLIGHT: In this work, we develop a framework to fuse both the single-view feature tracks and multi-view detected part locations to significantly improve the detection, localization and reconstruction of moving vehicles, even in the presence of strong occlusions.

198, **TITLE:** Deep Material-Aware Cross-Spectral Stereo Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhi_Deep_Material-Aware_Cross-Spectral_CVPR_2018_paper.html
AUTHORS: Tiancheng Zhi, Bernardo R. Pires, Martial Hebert, Srinivasa G. Narasimhan
HIGHLIGHT: We develop a novel deep learning framework to simultaneously transform images across spectral bands and estimate disparity.

199, **TITLE:** Augmenting Crowd-Sourced 3D Reconstructions Using Semantic Detections
http://openaccess.thecvf.com/content_cvpr_2018/html/Price_Augmenting_Crowd-Sourced_3D_CVPR_2018_paper.html
AUTHORS: True Price, Johannes L. Schönberger, Zhen Wei, Marc Pollefeys, Jan-Michael Frahm
HIGHLIGHT: We propose a method to jointly address these remaining open problems of SfM.

- 200, TITLE: Matryoshka Networks: Predicting 3D Geometry via Nested Shape Layers
http://openaccess.thecvf.com/content_cvpr_2018/html/Richter_Matryoshka_Networks_Predicting_CVPR_2018_paper.html
AUTHORS: Stephan R. Richter, Stefan Roth
HIGHLIGHT: In this paper, we develop novel, efficient 2D encodings for 3D geometry, which enable reconstructing full 3D shapes from a single image at high resolution.
- 201, TITLE: Triplet-Center Loss for Multi-View 3D Object Retrieval
http://openaccess.thecvf.com/content_cvpr_2018/html/He_Triplet-Center_Loss_for_CVPR_2018_paper.html
AUTHORS: Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, Xiang Bai
HIGHLIGHT: In the paper, we study variants of deep metric learning losses for 3D object retrieval, which did not receive enough attention from this area.
- 202, TITLE: Learning 3D Shape Completion From Laser Scan Data With Weak Supervision
http://openaccess.thecvf.com/content_cvpr_2018/html/Stutz_Learning_3D_Shape_CVPR_2018_paper.html
AUTHORS: David Stutz, Andreas Geiger
HIGHLIGHT: In this work, we propose a weakly-supervised learning-based approach to 3D shape completion which neither requires slow optimization nor direct supervision.
- 203, TITLE: End-to-End Learning of Keypoint Detector and Descriptor for Pose Invariant 3D Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Georgakis_End-to-End_Learning_of_CVPR_2018_paper.html
AUTHORS: Georgios Georgakis, Srikrishna Karanam, Ziyang Wu, Jan Ernst, Jana Košecá
HIGHLIGHT: This paper proposes an end-to-end learning framework for keypoint detection and its representation (descriptor) for 3D depth maps or 3D scans, where the two can be jointly optimized towards task-specific objectives without a need for separate annotations.
- 204, TITLE: ICE-BA: Incremental, Consistent and Efficient Bundle Adjustment for Visual-Inertial SLAM
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_ICE-BA_Incremental_Consistent_CVPR_2018_paper.html
AUTHORS: Haomin Liu, Mingyu Chen, Guofeng Zhang, Hujun Bao, Yingze Bao
HIGHLIGHT: In this work, we renovate the numerical solver for VI-SLAM.
- 205, TITLE: GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose
http://openaccess.thecvf.com/content_cvpr_2018/html/Yin_GeoNet_Unsupervised_Learning_CVPR_2018_paper.html
AUTHORS: Zhichao Yin, Jianping Shi
HIGHLIGHT: We propose GeoNet, a jointly unsupervised learning framework for monocular depth, optical flow and ego-motion estimation from videos.
- 206, TITLE: Radially-Distorted Conjugate Translations
http://openaccess.thecvf.com/content_cvpr_2018/html/Pritts_Radially-Distorted_Conjugate_Translations_CVPR_2018_paper.html
AUTHORS: James Pritts, Zuzana Kukelova, Viktor Larsson, Ondřej Chum
HIGHLIGHT: This paper introduces the first minimal solvers that jointly solve for affine-rectification and radial lens distortion from coplanar repeated patterns.
- 207, TITLE: Deep Ordinal Regression Network for Monocular Depth Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Fu_Deep_Ordinal_Regression_CVPR_2018_paper.html
AUTHORS: Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Dacheng Tao
HIGHLIGHT: To eliminate or at least largely reduce these problems, we introduce a spacing-increasing discretization (SID) strategy to discretize depth and recast depth network learning as an ordinal regression problem.
- 208, TITLE: Analytical Modeling of Vanishing Points and Curves in Catadioptric Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Miraldo_Analytical_Modeling_of_CVPR_2018_paper.html
AUTHORS: Pedro Miraldo, Francisco Eiras, Srikumar Ramalingam
HIGHLIGHT: In this paper, we derive parametric equations for vanishing points and vanishing curves using the calibration parameters, mirror shape coefficients, and direction vectors of parallel lines in 3D space.
- 209, TITLE: Learning Depth From Monocular Videos Using Direct Methods
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Learning_Depth_From_CVPR_2018_paper.html
AUTHORS: Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, Simon Lucey
HIGHLIGHT: Inspired by recent advances in direct visual odometry (DVO), we argue that the depth CNN predictor can be learned without a pose CNN predictor.

- 210, TITLE: Saliency Guided Depth Calibration for Perceptually Optimized Compressive Light Field 3D Display
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Saliency_Guided_Depth_CVPR_2018_paper.html
AUTHORS: Shizheng Wang, Wenjuan Liao, Phil Surman, Zhigang Tu, Yuanjin Zheng, Junsong Yuan
HIGHLIGHT: Considering this disadvantage, our paper incorporates a saliency guided depth optimization over a limited display range to calibrate the displayed depth and present the maximum area of saliency region for multi-layer light field display.
- 211, TITLE: MegaDepth: Learning Single-View Depth Prediction From Internet Photos
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_MegaDepth_Learning_Single-View_CVPR_2018_paper.html
AUTHORS: Zhengqi Li, Noah Snavely
HIGHLIGHT: We propose to use multi-view Internet photo collections, a virtually unlimited data source, to generate training data via modern structure-from-motion and multi-view stereo (MVS) methods, and present a large depth dataset called MegaDepth based on this idea.
- 212, TITLE: LayoutNet: Reconstructing the 3D Room Layout From a Single RGB Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Zou_LayoutNet_Reconstructing_the_CVPR_2018_paper.html
AUTHORS: Chuhan Zou, Alex Colburn, Qi Shan, Derek Hoiem
HIGHLIGHT: We propose an algorithm to predict room layout from a single image that generalizes across panoramas and perspective images, cuboid layouts and more general layouts (e.g. "L"-shape room).
- 213, TITLE: CBMV: A Coalesced Bidirectional Matching Volume for Disparity Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Batsos_CBMV_A_Coalesced_CVPR_2018_paper.html
AUTHORS: Konstantinos Batsos, Changjiang Cai, Philippos Mordohai
HIGHLIGHT: In this paper, we generate a matching volume leveraging both data with ground truth and conventional wisdom.
- 214, TITLE: Zoom and Learn: Generalizing Deep Stereo Matching to Novel Domains
http://openaccess.thecvf.com/content_cvpr_2018/html/Pang_Zoom_and_Learn_CVPR_2018_paper.html
AUTHORS: Jiahao Pang, Wenxiu Sun, Chengxi Yang, Jimmy Ren, Ruichao Xiao, Jin Zeng, Liang Lin
HIGHLIGHT: In this work, we propose a self-adaptation approach for CNN training, utilizing both synthetic training data (with ground-truth disparities) and stereo pairs in the new domain (without ground-truths).
- 215, TITLE: Exploring Disentangled Feature Representation Beyond Face Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Exploring_Disentangled_Feature_CVPR_2018_paper.html
AUTHORS: Yu Liu, Fangyin Wei, Jing Shao, Lu Sheng, Junjie Yan, Xiaogang Wang
HIGHLIGHT: This paper proposes learning disentangled but complementary face features with a minimal supervision by face identification.
- 216, TITLE: Learning Facial Action Units From Web Images With Scalable Weakly Supervised Clustering
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Learning_Facial_Action_CVPR_2018_paper.html
AUTHORS: Kaili Zhao, Wen-Sheng Chu, Aleix M. Martinez
HIGHLIGHT: We present a scalable weakly supervised clustering approach to learn facial action units (AUs) from large, freely available web images.
- 217, TITLE: Human Pose Estimation With Parsing Induced Learner
http://openaccess.thecvf.com/content_cvpr_2018/html/Nie_Human_Pose_Estimation_CVPR_2018_paper.html
AUTHORS: Xuecheng Nie, Jiashi Feng, Yiming Zuo, Shuicheng Yan
HIGHLIGHT: In this paper, we propose a novel Parsing Induced Learner to exploit parsing information to effectively assist pose estimation by learning to fast adapt the base pose estimation model.
- 218, TITLE: Multi-Level Factorisation Net for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Chang_Multi-Level_Factorisation_Net_CVPR_2018_paper.html
AUTHORS: Xiaobin Chang, Timothy M. Hospedales, Tao Xiang
HIGHLIGHT: We propose Multi-Level Factorisation Net (MLFN), a novel network architecture that factorises the visual appearance of a person into latent discriminative factors at multiple semantic levels without manual annotation.
- 219, TITLE: Attention-Aware Compositional Network for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Attention-Aware_Compositional_Network_CVPR_2018_paper.html
AUTHORS: Jing Xu, Rui Zhao, Feng Zhu, Huaming Wang, Wanli Ouyang
HIGHLIGHT: In this work, we introduce a novel framework called Attention-Aware Compositional Network (AACN) for person ReID.

- 220, TITLE: Look at Boundary: A Boundary-Aware Face Alignment Algorithm
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Look_at_Boundary_CVPR_2018_paper.html
AUTHORS: Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, Qiang Zhou
HIGHLIGHT: We present a novel boundary-aware face alignment algorithm by utilising boundary lines as the geometric structure of a human face to help facial landmark localisation.
- 221, TITLE: Demo2Vec: Reasoning Object Affordances From Online Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Fang_Demo2Vec_Reasoning_Object_CVPR_2018_paper.html
AUTHORS: Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, Joseph J. Lim
HIGHLIGHT: In this paper, we consider the problem of reasoning object affordances through the feature embedding of demonstration videos.
We introduce the Online Product Review dataset for Affordance (OPRA) by collecting and labeling diverse YouTube product review videos.
- 222, TITLE: Monocular 3D Pose and Shape Estimation of Multiple People in Natural Scenes - The Importance of Multiple Scene Constraints
http://openaccess.thecvf.com/content_cvpr_2018/html/Zanfir_Monocular_3D_Pose_CVPR_2018_paper.html
AUTHORS: Andrei Zanfir, Elisabeta Marinoiu, Cristian Sminchisescu
HIGHLIGHT: In this paper, we leverage state-of-the-art deep multi-task neural networks and parametric human and scene modeling, towards a fully automatic monocular visual sensing system for multiple interacting people, which (i) infers the 2d and 3d pose and shape of multiple people from a single image, relying on detailed semantic representations at both model and image level, to guide a combined optimization with feedforward and feedback components, (ii) automatically integrates scene constraints including ground plane support and simultaneous volume occupancy by multiple people, and (iii) extends the single image model to video by optimally solving the temporal person assignment problem and imposing coherent temporal pose and motion reconstructions while preserving image alignment fidelity.
- 223, TITLE: 3D Human Sensing, Action and Emotion Recognition in Robot Assisted Therapy of Children With Autism
http://openaccess.thecvf.com/content_cvpr_2018/html/Marinoiu_3D_Human_Sensing_CVPR_2018_paper.html
AUTHORS: Elisabeta Marinoiu, Mihai Zanfir, Vlad Olaru, Cristian Sminchisescu
HIGHLIGHT: We introduce new, fine-grained action and emotion recognition tasks defined on non-staged videos, recorded during robot-assisted therapy sessions of children with autism.
- 224, TITLE: Facial Expression Recognition by De-Expression Residue Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Facial_Expression_Recognition_CVPR_2018_paper.html
AUTHORS: Huiyuan Yang, Umur Ciftci, Lijun Yin
HIGHLIGHT: In this paper, we propose to recognize facial expressions by extracting information of the expressive component through a de-expression learning procedure, called De-expression Residue Learning (DeRL).
- 225, TITLE: A Causal And-Or Graph Model for Visibility Fluent Reasoning in Tracking Interacting Objects
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_A_Causal_And-Or_CVPR_2018_paper.html
AUTHORS: Yuanlu Xu, Lei Qin, Xiaobai Liu, Jianwen Xie, Song-Chun Zhu
HIGHLIGHT: In this work, we consider the visibility status of a subject as a fluent variable, whose change is mostly attributed to the subject's interaction with the surrounding, e.g., crossing behind another object, entering a building, or getting into a vehicle, etc.
- 226, TITLE: Weakly Supervised Facial Action Unit Recognition Through Adversarial Training
http://openaccess.thecvf.com/content_cvpr_2018/html/Peng_Weakly_Supervised_Facial_CVPR_2018_paper.html
AUTHORS: Guozhu Peng, Shangfei Wang
HIGHLIGHT: In this paper, we propose a novel AU recognition method that learns AU classifiers from domain knowledge and expression-annotated facial images through adversarial training.
- 227, TITLE: Non-Linear Temporal Subspace Representations for Activity Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Cherian_Non-Linear_Temporal_Subspace_CVPR_2018_paper.html
AUTHORS: Anoop Cherian, Suvrit Sra, Stephen Gould, Richard Hartley
HIGHLIGHT: As these features are often non-linear, we propose a novel pooling method, kernelized rank pooling, that represents a given sequence compactly as the pre-image of the parameters of a hyperplane in a reproducing kernel Hilbert space, projections of data onto which captures their temporal order.
- 228, TITLE: Towards Pose Invariant Face Recognition in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Towards_Pose_Invariant_CVPR_2018_paper.html
AUTHORS: Jian Zhao, Yu Cheng, Yan Xu, Lin Xiong, Jianshu Li, Fang Zhao, Karlekar Jayashree, Sugiri Pranata, Shengmei Shen, Junliang Xing, Shuicheng Yan, Jiashi Feng

- HIGHLIGHT:** To this end, we propose a Pose Invariant Model (PIM) for face recognition in the wild, with three distinct novelties.
- 229, **TITLE:** Unifying Identification and Context Learning for Person Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Unifying_Identification_and_CVPR_2018_paper.html
AUTHORS: Qingqiu Huang, Yu Xiong, Dahua Lin
HIGHLIGHT: In this work, we aim to move beyond such limitations and propose a new framework to leverage context for person recognition.
- 230, **TITLE:** Jointly Optimize Data Augmentation and Network Training: Adversarial Data Augmentation in Human Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Peng_Jointly_Optimize_Data_CVPR_2018_paper.html
AUTHORS: Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S. Feris, Dimitris Metaxas
HIGHLIGHT: We propose adversarial data augmentation to address this limitation.
- 231, **TITLE:** Wing Loss for Robust Facial Landmark Localisation With Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Feng_Wing_Loss_for_CVPR_2018_paper.html
AUTHORS: Zhen-Hua Feng, Josef Kittler, Muhammad Awaiz, Patrik Huber, Xiao-Jun Wu
HIGHLIGHT: We present a new loss function, namely Wing loss, for robust facial landmark localisation with Convolutional Neural Networks (CNNs).
- 232, **TITLE:** Multiple Granularity Group Interaction Prediction
http://openaccess.thecvf.com/content_cvpr_2018/html/Yao_Multiple_Granularity_Group_CVPR_2018_paper.html
AUTHORS: Taiping Yao, Minsi Wang, Bingbing Ni, Huawei Wei, Xiaokang Yang
HIGHLIGHT: In contrast, in this work, we propose a multi-granularity interaction prediction network which integrates both global motion and detailed local action.
- 233, **TITLE:** Social GAN: Socially Acceptable Trajectories With Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Gupta_Social_GAN_Socially_CVPR_2018_paper.html
AUTHORS: Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, Alexandre Alahi
HIGHLIGHT: We tackle this problem by combining tools from sequence prediction and generative adversarial networks: a recurrent sequence-to-sequence model observes motion histories and predicts future behavior, using a novel pooling mechanism to aggregate information across people.
- 234, **TITLE:** Deep Group-Shuffling Random Walk for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Deep_Group-Shuffling_Random_CVPR_2018_paper.html
AUTHORS: Yantao Shen, Hongsheng Li, Tong Xiao, Shuai Yi, Dapeng Chen, Xiaogang Wang
HIGHLIGHT: In this paper, we propose a novel group-shuffling random walk network for fully utilizing the affinity information between gallery images in both the training and testing processes.
- 235, **TITLE:** Transferable Joint Attribute-Identity Deep Learning for Unsupervised Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Transferable_Joint_Attribute-Identity_CVPR_2018_paper.html
AUTHORS: Jingya Wang, Xiatian Zhu, Shaogang Gong, Wei Li
HIGHLIGHT: Specifically, we introduce an Transferable Joint Attribute-Identity Deep Learning (TJ-AIDL) for simultaneously learning an attribute-semantic and identitydiscriminative feature representation space transferrable to any new (unseen) target domain for re-id tasks without the need for collecting new labelled training data from the target domain (i.e. unsupervised learning in the target domain).
- 236, **TITLE:** Harmonious Attention Network for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Harmonious_Attention_Network_CVPR_2018_paper.html
AUTHORS: Wei Li, Xiatian Zhu, Shaogang Gong
HIGHLIGHT: In this work, we show the advantages of jointly learning attention selection and feature representation in a Convolutional Neural Network (CNN) by maximising the complementary information of different levels of visual attention subject to re-id discriminative learning constraints.
- 237, **TITLE:** Real-Time Rotation-Invariant Face Detection With Progressive Calibration Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Shi_Real-Time_Rotation-Invariant_Face_CVPR_2018_paper.html
AUTHORS: Xuepeng Shi, Shiguang Shan, Meina Kan, Shuzhe Wu, Xilin Chen
HIGHLIGHT: To address this problem more efficiently, we propose Progressive Calibration Networks (PCN) to perform rotation-invariant face detection in a coarse-to-fine manner.

- 238, TITLE: Deep Regression Forests for Age Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Deep_Regression_Forests_CVPR_2018_paper.html
AUTHORS: Wei Shen, Yilu Guo, Yan Wang, Kai Zhao, Bo Wang, Alan L. Yuille
HIGHLIGHT: In this paper, we propose Deep Regression Forests (DRFs), an end-to-end model, for age estimation.
- 239, TITLE: Weakly-Supervised Deep Convolutional Neural Network Learning for Facial Action Unit Intensity Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Weakly-Supervised_Deep_Convolutional_CVPR_2018_paper.html
AUTHORS: Yong Zhang, Weiming Dong, Bao-Gang Hu, Qiang Ji
HIGHLIGHT: Recent works have introduced deep neural networks for AU intensity estimation, but they require a large amount of intensity annotations.
- 240, TITLE: Memory Based Online Learning of Deep Representations From Video Streams
http://openaccess.thecvf.com/content_cvpr_2018/html/Pernici_Memory_Based_Online_CVPR_2018_paper.html
AUTHORS: Federico Pernici, Federico Bartoli, Matteo Bruni, Alberto Del Bimbo
HIGHLIGHT: We present a novel online unsupervised method for face identity learning from video streams.
- 241, TITLE: Efficient and Deep Person Re-Identification Using Multi-Level Similarity
http://openaccess.thecvf.com/content_cvpr_2018/html/Guo_Efficient_and_Deep_CVPR_2018_paper.html
AUTHORS: Yiluan Guo, Ngai-Man Cheung
HIGHLIGHT: In this work, we propose an efficient, end-to-end fully convolutional Siamese network that computes the similarities at multiple levels.
- 242, TITLE: Multi-Level Fusion Based 3D Object Detection From Monocular Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Multi-Level_Fusion_Based_CVPR_2018_paper.html
AUTHORS: Bin Xu, Zhenzhong Chen
HIGHLIGHT: In this paper, we present an end-to-end deep learning based framework for 3D object detection from a single monocular image.
- 243, TITLE: A Perceptual Measure for Deep Single Image Camera Calibration
http://openaccess.thecvf.com/content_cvpr_2018/html/Hold-Geoffroy_A_Perceptual_Measure_CVPR_2018_paper.html
AUTHORS: Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, Jean-François Lalonde
HIGHLIGHT: We propose inferring directly camera calibration parameters from a single image using a deep convolutional neural network.
- 244, TITLE: Learning to Generate Time-Lapse Videos Using Multi-Stage Dynamic Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Xiong_Learning_to_Generate_CVPR_2018_paper.html
AUTHORS: Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, Jiebo Luo
HIGHLIGHT: We address this problem by presenting a generative adversarial network (GAN) based two-stage approach to generating realistic time-lapse videos of high resolution.
We build a large scale time-lapse dataset, and test our approach on this new dataset.
- 245, TITLE: Document Enhancement Using Visibility Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Kligler_Document_Enhancement_Using_CVPR_2018_paper.html
AUTHORS: Netanel Kligler, Sagi Katz, Ayellet Tal
HIGHLIGHT: Rather than proposing a new algorithm for a specific problem, we introduce a novel general approach.
- 246, TITLE: A Weighted Sparse Sampling and Smoothing Frame Transition Approach for Semantic Fast-Forward First-Person Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Silva_A_Weighted_Sparse_CVPR_2018_paper.html
AUTHORS: Michel Silva, Washington Ramos, João Ferreira, Felipe Chamone, Mario Campos, Erickson R. Nascimento
HIGHLIGHT: In this work, we address the problem of creating smooth fast-forward videos without losing the relevant content.
- 247, TITLE: Context Contrasted Feature and Gated Multi-Scale Aggregation for Scene Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Ding_Context_Contrasted_Feature_CVPR_2018_paper.html
AUTHORS: Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, Gang Wang
HIGHLIGHT: Without bells and whistles, the proposed approach achieves the state-of-the-arts consistently on the three popular scene segmentation datasets, Pascal Context, SUN-RGBD and COCO Stuff.

- 248, TITLE: Deep Layer Aggregation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Deep_Layer_Aggregation_CVPR_2018_paper.html
AUTHORS: Fisher Yu, Dequan Wang, Evan Shelhamer, Trevor Darrell
HIGHLIGHT: Architectural efforts are exploring many dimensions for network backbones, designing deeper or wider architectures, but how to best aggregate layers and blocks across a network deserves further attention.
- 249, TITLE: Convolutional Neural Networks With Alternately Updated Clique
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Convolutional_Neural_Networks_CVPR_2018_paper.html
AUTHORS: Yibo Yang, Zhisheng Zhong, Tiancheng Shen, Zhouchen Lin
HIGHLIGHT: Here we propose a new convolutional neural network architecture with alternately updated clique (CliqueNet).
- 250, TITLE: Practical Block-Wise Neural Network Architecture Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhong_Practical_Block-Wise_Neural_CVPR_2018_paper.html
AUTHORS: Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, Cheng-Lin Liu
HIGHLIGHT: In this paper, we provide a block-wise network generation pipeline called BlockQNN which automatically builds high-performance networks using the Q-Learning paradigm with epsilon-greedy exploration strategy.
- 251, TITLE: xUnit: Learning a Spatial Activation Function for Efficient Image Restoration
http://openaccess.thecvf.com/content_cvpr_2018/html/Kligvasser_xUnit_Learning_a_CVPR_2018_paper.html
AUTHORS: Idan Kligvasser, Tamar Rott Shaham, Tomer Michaeli
HIGHLIGHT: In this paper, we propose a new activation unit, which is particularly suitable for image restoration problems.
- 252, TITLE: Crafting a Toolchain for Image Restoration by Deep Reinforcement Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Crafting_a_Toolchain_CVPR_2018_paper.html
AUTHORS: Ke Yu, Chao Dong, Liang Lin, Chen Change Loy
HIGHLIGHT: We investigate a novel approach for image restoration by reinforcement learning.
- 253, TITLE: Deformation Aware Image Compression
http://openaccess.thecvf.com/content_cvpr_2018/html/Shaham_Deformation_Aware_Image_CVPR_2018_paper.html
AUTHORS: Tamar Rott Shaham, Tomer Michaeli
HIGHLIGHT: Lossy compression algorithms aim to compactly encode images in a way which enables to restore them with minimal error.
- 254, TITLE: Distributable Consistent Multi-Object Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Distributable_Consistent_Multi-Object_CVPR_2018_paper.html
AUTHORS: Nan Hu, Qixing Huang, Boris Thibert, Leonidas J. Guibas
HIGHLIGHT: In this paper we propose an optimization-based framework to multiple object matching.
- 255, TITLE: Residual Dense Network for Image Super-Resolution
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Residual_Dense_Network_CVPR_2018_paper.html
AUTHORS: Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, Yun Fu
HIGHLIGHT: In this paper, we propose dense feature fusion (DFF) for image super-resolution (SR).
- 256, TITLE: Attentive Generative Adversarial Network for Raindrop Removal From a Single Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Qian_Attentive_Generative_Adversarial_CVPR_2018_paper.html
AUTHORS: Rui Qian, Robby T. Tan, Wenhan Yang, Jiajun Su, Jiaying Liu
HIGHLIGHT: In this paper, we address the problem by visually removing raindrops, and thus transforming a raindrop degraded image into a clean one.
- 257, TITLE: FSRNet: End-to-End Learning Face Super-Resolution With Facial Priors
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_FSRNet_End-to-End_Learning_CVPR_2018_paper.html
AUTHORS: Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, Jian Yang
HIGHLIGHT: We present a novel deep end-to-end trainable Face Super-Resolution Network (FSRNet), which makes use of the geometry prior, i.e., facial landmark heatmaps and parsing maps, to superresolve very low-resolution (LR) face images without wellaligned requirement.
- 258, TITLE: Burst Denoising With Kernel Prediction Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Mildenhall_Burst_Denoising_With_CVPR_2018_paper.html
AUTHORS: Ben Mildenhall, Jonathan T. Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, Robert Carroll
HIGHLIGHT: We present a technique for jointly denoising bursts of images taken from a handheld camera.

- 259, TITLE: Unsupervised Sparse Dirichlet-Net for Hyperspectral Image Super-Resolution
http://openaccess.thecvf.com/content_cvpr_2018/html/Qu_Unsupervised_Sparse_Dirichlet-Net_CVPR_2018_paper.html
AUTHORS: Ying Qu, Hairong Qi, Chiman Kwan
HIGHLIGHT: This paper focuses on hyperspectral image super-resolution (HSI-SR), where a hyperspectral image (HSI) with low spatial resolution (LR) but high spectral resolution is fused with a multispectral image (MSI) with high spatial resolution (HR) but low spectral resolution to obtain HR HSI.
- 260, TITLE: Dynamic Scene Deblurring Using Spatially Variant Recurrent Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Dynamic_Scene_Deblurring_CVPR_2018_paper.html
AUTHORS: Jiawei Zhang, Jinshan Pan, Jimmy Ren, Yibing Song, Linchao Bao, Rynson W.H. Lau, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a novel spatially variant neural network to address the problem.
- 261, TITLE: SPLATNet: Sparse Lattice Networks for Point Cloud Processing
http://openaccess.thecvf.com/content_cvpr_2018/html/Su_SPLATNet_Sparse_Lattice_CVPR_2018_paper.html
AUTHORS: Hang Su, Varun Jampani, Deqing Sun, Subhansu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, Jan Kautz
HIGHLIGHT: We present a network architecture for processing point clouds that directly operates on a collection of points represented as a sparse set of samples in a high-dimensional lattice.
- 262, TITLE: Surface Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Kostrikov_Surface_Networks_CVPR_2018_paper.html
AUTHORS: Ilya Kostrikov, Zhongshi Jiang, Daniele Panozzo, Denis Zorin, Joan Bruna
HIGHLIGHT: To overcome this limitation, we propose several upgrades to GNNs to leverage extrinsic differential geometry properties of three-dimensional surfaces, increasing its modeling power.
- 263, TITLE: Self-Supervised Multi-Level Face Model Learning for Monocular Reconstruction at Over 250 Hz
http://openaccess.thecvf.com/content_cvpr_2018/html/Tewari_Self-Supervised_Multi-Level_Face_CVPR_2018_paper.html
AUTHORS: Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, Christian Theobalt
HIGHLIGHT: To alleviate this problem, we present the first approach that jointly learns 1) a regressor for face shape, expression, reflectance and illumination on the basis of 2) a concurrently learned parametric face model.
- 264, TITLE: CodeSLAM — Learning a Compact, Optimisable Representation for Dense Visual SLAM
http://openaccess.thecvf.com/content_cvpr_2018/html/Bloesch_CodeSLAM_-_Learning_CVPR_2018_paper.html
AUTHORS: Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, Andrew J. Davison
HIGHLIGHT: We present a new compact but dense representation of scene geometry which is conditioned on the intensity data from a single image and generated from a code consisting of a small number of parameters.
- 265, TITLE: SGPN: Similarity Group Proposal Network for 3D Point Cloud Instance Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_SGPN_Similarity_Group_CVPR_2018_paper.html
AUTHORS: Weiyue Wang, Ronald Yu, Qiangui Huang, Ulrich Neumann
HIGHLIGHT: We introduce Similarity Group Proposal Network (SGPN), a simple and intuitive deep learning framework for 3D object instance segmentation on point clouds.
- 266, TITLE: PlaneNet: Piece-Wise Planar Reconstruction From a Single RGB Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_PlaneNet_Piece-Wise_Planar_CVPR_2018_paper.html
AUTHORS: Chen Liu, Jiméi Yang, Duygu Ceylan, Ersin Yumer, Yasutaka Furukawa
HIGHLIGHT: This paper proposes a deep neural network (DNN) for piece-wise planar depthmap reconstruction from a single RGB image.
- 267, TITLE: Deep Parametric Continuous Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Deep_Parametric_Continuous_CVPR_2018_paper.html
AUTHORS: Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, Raquel Urtasun
HIGHLIGHT: In this paper we propose Parametric Continuous Convolution, a new learnable operator that operates over non-grid structured data.
- 268, TITLE: FeaStNet: Feature-Steered Graph Convolutions for 3D Shape Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Verma_FeaStNet_Feature-Steered_Graph_CVPR_2018_paper.html
AUTHORS: Nitika Verma, Edmond Boyer, Jakob Verbeek

HIGHLIGHT: To address this problem, we propose a novel graph-convolution operator to establish correspondences between filter weights and graph neighborhoods with arbitrary connectivity.

269, **TITLE:** Image Collection Pop-Up: 3D Reconstruction and Clustering of Rigid and Non-Rigid Categories
http://openaccess.thecvf.com/content_cvpr_2018/html/Agudo_Image_Collection_Pop-Up_CVPR_2018_paper.html
AUTHORS: Antonio Agudo, Melcior Pijoan, Francesc Moreno-Noguer
HIGHLIGHT: This paper introduces an approach to simultaneously estimate 3D shape, camera pose, and object and type of deformation clustering, from partial 2D annotations in a multi-instance collection of images.

270, **TITLE:** Geometry-Aware Learning of Maps for Camera Localization
http://openaccess.thecvf.com/content_cvpr_2018/html/Brahmbhatt_Geometry-Aware_Learning_of_CVPR_2018_paper.html
AUTHORS: Samarth Brahmbhatt, Jinwei Gu, Kihwan Kim, James Hays, Jan Kautz
HIGHLIGHT: We propose to represent maps as a deep neural net called MapNet, which enables learning a data-driven map representation.

271, **TITLE:** Recurrent Slice Networks for 3D Segmentation of Point Clouds
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Recurrent_Slice_Networks_CVPR_2018_paper.html
AUTHORS: Qiangui Huang, Weiyue Wang, Ulrich Neumann
HIGHLIGHT: This work presents a novel 3D segmentation framework, RSNet, to efficiently model local structures in point clouds.

272, **TITLE:** Depth-Based 3D Hand Pose Estimation: From Current Achievements to Future Goals
http://openaccess.thecvf.com/content_cvpr_2018/html/Yuan_Depth-Based_3D_Hand_CVPR_2018_paper.html
AUTHORS: Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Lihao Ge, Junsong Yuan, Xinghao Chen, Guijin Wang, Fan Yang, Kai Akiyama, Yang Wu, Qingfu Wan, Meysam Madadi, Sergio Escalera, Shile Li, Dongheui Lee, Iason Oikonomidis, Antonis Argyros, Tae-Kyun Kim
HIGHLIGHT: In this paper, we strive to answer two questions: What is the current state of 3D hand pose estimation from depth images?

273, **TITLE:** SobolevFusion: 3D Reconstruction of Scenes Undergoing Free Non-Rigid Motion
http://openaccess.thecvf.com/content_cvpr_2018/html/Slavcheva_SobolevFusion_3D_Reconstruction_CVPR_2018_paper.html
AUTHORS: Miroslava Slavcheva, Maximilian Baust, Slobodan Ilic
HIGHLIGHT: We present a system that builds 3D models of non-rigidly moving surfaces from scratch in real time using a single RGB-D stream.

274, **TITLE:** AdaDepth: Unsupervised Content Congruent Adaptation for Depth Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Kundu_AdaDepth_Unsupervised_content_cvpr_2018_paper.html
AUTHORS: Jogendra Nath Kundu, Phani Krishna Uppala, Anuj Pahuja, R. Venkatesh Babu
HIGHLIGHT: In this work, we propose AdaDepth - an unsupervised domain adaptation strategy for the pixel-wise regression task of monocular depth estimation.

275, **TITLE:** Learning to Find Good Correspondences
http://openaccess.thecvf.com/content_cvpr_2018/html/Yi_Learning_to_Find_CVPR_2018_paper.html
AUTHORS: Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: We develop a deep architecture to learn to find good correspondences for wide-baseline stereo.

276, **TITLE:** OATM: Occlusion Aware Template Matching by Consensus Set Maximization
http://openaccess.thecvf.com/content_cvpr_2018/html/Korman_OATM_Occlusion_Aware_CVPR_2018_paper.html
AUTHORS: Simon Korman, Mark Milam, Stefano Soatto
HIGHLIGHT: We present a novel approach to template matching that is efficient, can handle partial occlusions, and comes with provable performance guarantees.

277, **TITLE:** Deep Learning of Graph Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Zanfir_Deep_Learning_of_CVPR_2018_paper.html
AUTHORS: Andrei Zanfir, Cristian Sminchisescu
HIGHLIGHT: We present an end-to-end model that makes it possible to learn all parameters of the graph matching process, including the unary and pairwise node neighborhoods, represented as deep feature extraction hierarchies.

278, **TITLE:** Unsupervised Discovery of Object Landmarks as Structural Representations

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Unsupervised_Discovery_of_CVPR_2018_paper.html

AUTHORS: Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, Honglak Lee

HIGHLIGHT: We propose an autoencoding formulation to discover landmarks as explicit structural representations.

279, TITLE: Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference

http://openaccess.thecvf.com/content_cvpr_2018/html/Jacob_Quantization_and_Training_CVPR_2018_paper.html

AUTHORS: Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, Dmitry Kalenichenko

HIGHLIGHT: We propose a quantization scheme along with a co-designed training procedure allowing inference to be carried out using integer-only arithmetic while preserving an end-to-end model accuracy that is close to floating-point inference.

280, TITLE: Lean Multiclass Crowdsourcing

http://openaccess.thecvf.com/content_cvpr_2018/html/Van_Horn_Lean_Multiclass_Crowdsourcing_CVPR_2018_paper.html

AUTHORS: Grant Van Horn, Steve Branson, Scott Loarie, Serge Belongie, Pietro Perona

HIGHLIGHT: We introduce a method for efficiently crowdsourcing multiclass annotations in challenging, real world image datasets.

281, TITLE: Partial Transfer Learning With Selective Adversarial Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Partial_Transfer_Learning_CVPR_2018_paper.html

AUTHORS: Zhangjie Cao, Mingsheng Long, Jianmin Wang, Michael I. Jordan

HIGHLIGHT: This paper introduces partial transfer learning, which relaxes the shared label space assumption to that the target label space is only a subspace of the source label space.

282, TITLE: Self-Supervised Feature Learning by Learning to Spot Artifacts

http://openaccess.thecvf.com/content_cvpr_2018/html/Jenni_Self-Supervised_Feature_Learning_CVPR_2018_paper.html

AUTHORS: Simon Jenni, Paolo Favaro

HIGHLIGHT: We introduce a novel self-supervised learning method based on adversarial training.

283, TITLE: LDMNet: Low Dimensional Manifold Regularized Neural Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_LDMNet_Low_Dimensional_CVPR_2018_paper.html

AUTHORS: Wei Zhu, Qiang Qiu, Jiayi Huang, Robert Calderbank, Guillermo Sapiro, Ingrid Daubechies

HIGHLIGHT: To resolve this, we propose the Low-Dimensional- Manifold-regularized neural Network (LDMNet), which incorporates a feature regularization method that focuses on the geometry of both the input data and the output features.

284, TITLE: CondenseNet: An Efficient DenseNet Using Learned Group Convolutions

http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_CondenseNet_An_Efficient_CVPR_2018_paper.html

AUTHORS: Gao Huang, Shichen Liu, Laurens van der Maaten, Kilian Q. Weinberger

HIGHLIGHT: In this paper we develop CondenseNet, a novel network architecture with unprecedented efficiency.

285, TITLE: Learning Deep Descriptors With Scale-Aware Triplet Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Keller_Learning_Deep_Descriptors_CVPR_2018_paper.html

AUTHORS: Michel Keller, Zetao Chen, Fabiola Maffra, Patrik Schmuck, Margarita Chli

HIGHLIGHT: Based on this analysis, we introduce mixed-context losses and scale-aware sampling, two methods that when combined enable networks to learn consistently scaled descriptors for the first time.

286, TITLE: Decoupled Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Decoupled_Networks_CVPR_2018_paper.html

AUTHORS: Weiyang Liu, Zhen Liu, Zhiding Yu, Bo Dai, Rongmei Lin, Yisen Wang, James M. Rehg, Le Song

HIGHLIGHT: Inspired by the observation that CNN-learned features are naturally decoupled with the norm of features corresponding to the intra-class variation and the angle corresponding to the semantic difference, we propose a generic decoupled learning framework which models the intra-class variation and semantic difference independently.

287, TITLE: Deep Adversarial Metric Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Duan_Deep_Adversarial_Metric_CVPR_2018_paper.html

AUTHORS: Yueqi Duan, Wenzhao Zheng, Xudong Lin, Jiwen Lu, Jie Zhou

HIGHLIGHT: In this paper, we propose a deep adversarial metric learning (DAML) framework to generate synthetic hard negatives from the observed negative samples, which is widely applicable to supervised deep metric learning methods.

288, TITLE: PU-Net: Point Cloud Upsampling Network

http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_PU-Net_Point_Cloud_CVPR_2018_paper.html

- AUTHORS: Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, Pheng-Ann Heng
HIGHLIGHT: In this paper, we present a data-driven point cloud upsampling technique.
- 289, TITLE: Real-Time Monocular Depth Estimation Using Synthetic Data With Domain Adaptation via Image Style Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Atapour-Abarghouei_Real-Time_Monocular_Depth_CVPR_2018_paper.html
AUTHORS: Amir Atapour-Abarghouei, Toby P. Breckon
HIGHLIGHT: Monocular depth estimation using learning-based approaches has become promising in recent years.
- 290, TITLE: Learning for Disparity Estimation Through Feature Constancy
http://openaccess.thecvf.com/content_cvpr_2018/html/Liang_Learning_for_Disparity_CVPR_2018_paper.html
AUTHORS: Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, Jianfeng Zhang
HIGHLIGHT: In this paper, we propose a network architecture to incorporate all steps of stereo matching.
- 291, TITLE: DeepMVS: Learning Multi-View Stereopsis
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_DeepMVS_Learning_Multi-View_CVPR_2018_paper.html
AUTHORS: Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, Jia-Bin Huang
HIGHLIGHT: We present DeepMVS, a deep convolutional neural network (ConvNet) for multi-view stereo reconstruction.
- 292, TITLE: Self-Calibrating Polarising Radiometric Calibration
http://openaccess.thecvf.com/content_cvpr_2018/html/Teo_Self-Calibrating_Polarising_Radiometric_CVPR_2018_paper.html
AUTHORS: Daniel Teo, Boxin Shi, Yinqiang Zheng, Sai-Kit Yeung
HIGHLIGHT: We present a self-calibrating polarising radiometric calibration method.
- 293, TITLE: Coding Kendall's Shape Trajectories for 3D Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Tanfous_Coding_Kendalls_Shape_CVPR_2018_paper.html
AUTHORS: Amor Ben Tanfous, Hassen Drira, Boulbaba Ben Amor
HIGHLIGHT: This puts an additional constraint (i.e., non-linearity) in using conventional machine learning techniques for the purpose of classification, event detection, prediction, etc.
- 294, TITLE: Efficient, Sparse Representation of Manifold Distance Matrices for Classical Scaling
http://openaccess.thecvf.com/content_cvpr_2018/html/Turek_Efficient_Sparse_Representation_CVPR_2018_paper.html
AUTHORS: Javier S. Turek, Alexander G. Huth
HIGHLIGHT: In this paper we present a novel sparse method for efficiently representing geodesic distance matrices using biharmonic interpolation.
- 295, TITLE: Motion Segmentation by Exploiting Complementary Geometric Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Motion_Segmentation_by_CVPR_2018_paper.html
AUTHORS: Xun Xu, Loong Fah Cheong, Zhuwen Li
HIGHLIGHT: From these considerations, we propose a multi-view spectral clustering framework that synergistically combines multiple models together.
- 296, TITLE: Estimation of Camera Locations in Highly Corrupted Scenarios: All About That Base, No Shape Trouble
http://openaccess.thecvf.com/content_cvpr_2018/html/Shi_Estimation_of_Camera_CVPR_2018_paper.html
AUTHORS: Yunpeng Shi, Gilad Lerman
HIGHLIGHT: We propose a strategy for improving camera location estimation in structure from motion.
- 297, TITLE: 4D Human Body Correspondences From Panoramic Depth Maps
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_4D_Human_Body_CVPR_2018_paper.html
AUTHORS: Zhong Li, Minye Wu, Wangyiteng Zhou, Jingyi Yu
HIGHLIGHT: We present an end-to-end deep learning scheme to establish dense shape correspondences and subsequently compress the data.
- 298, TITLE: Reconstructing Thin Structures of Manifold Surfaces by Integrating Spatial Curves
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Reconstructing_Thin_Structures_CVPR_2018_paper.html
AUTHORS: Shiwei Li, Yao Yao, Tian Fang, Long Quan
HIGHLIGHT: In this paper, we address this problem by leveraging spatial curves.
- 299, TITLE: Multi-View Consistency as Supervisory Signal for Learning Shape and Pose Prediction

http://openaccess.thecvf.com/content_cvpr_2018/html/Tulsiani_Multi-View_Consistency_as_CVPR_2018_paper.html
AUTHORS: Shubham Tulsiani, Alexei A. Efros, Jitendra Malik
HIGHLIGHT: We present a framework for learning single-view shape and pose prediction without using direct supervision for either.

300, TITLE: Probabilistic Plant Modeling via Multi-View Image-to-Image Translation
http://openaccess.thecvf.com/content_cvpr_2018/html/Isokane_Probabilistic_Plant_Modeling_CVPR_2018_paper.html
AUTHORS: Takahiro Isokane, Fumio Okura, Ayaka Ide, Yasuyuki Matsushita, Yasushi Yagi
HIGHLIGHT: This paper describes a method for inferring three-dimensional (3D) plant branch structures that are hidden under leaves from multi-view observations.

301, TITLE: Deep Marching Cubes: Learning Explicit Surface Representations
http://openaccess.thecvf.com/content_cvpr_2018/html/Liao_Deep_Marching_Cubes_CVPR_2018_paper.html
AUTHORS: Yiyi Liao, Simon Donné, Andreas Geiger
HIGHLIGHT: In this paper, we investigate the problem of end-to-end 3D surface prediction.
We further propose a set of loss functions which allow for training our model with sparse point supervision.

302, TITLE: Tags2Parts: Discovering Semantic Regions From Shape Tags
http://openaccess.thecvf.com/content_cvpr_2018/html/Muralikrishnan_Tags2Parts_Discovering_Semantic_CVPR_2018_paper.html
AUTHORS: Sanjeev Muralikrishnan, Vladimir G. Kim, Siddhartha Chaudhuri
HIGHLIGHT: We propose a novel method for discovering shape regions that strongly correlate with user-prescribed tags.

303, TITLE: Uncalibrated Photometric Stereo Under Natural Illumination
http://openaccess.thecvf.com/content_cvpr_2018/html/Mo_Uncalibrated_Photometric_Stereo_CVPR_2018_paper.html
AUTHORS: Zhipeng Mo, Boxin Shi, Feng Lu, Sai-Kit Yeung, Yasuyuki Matsushita
HIGHLIGHT: This paper presents a photometric stereo method that works with unknown natural illuminations without any calibration object.

304, TITLE: Robust Depth Estimation From Auto Bracketed Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Im_Robust_Depth_Estimation_CVPR_2018_paper.html
AUTHORS: Sunghoon Im, Hae-Gon Jeon, In So Kweon
HIGHLIGHT: To address the problem, we present a robust depth estimation method from a short burst shot with varied intensity (i.e., Auto Bracketing) or strong noise (i.e., High ISO).

305, TITLE: Free Supervision From Video Games
http://openaccess.thecvf.com/content_cvpr_2018/html/Krahenbuhl_Free_Supervision_From_CVPR_2018_paper.html
AUTHORS: Philipp Krähenbühl
HIGHLIGHT: We present an alternative, and show how ground truth labels for many vision tasks are easily extracted from video games in real time as we play them.
We collected a dataset of 220k training images, and 60k test images across 3 video games, and evaluate state of the art optical flow, depth estimation and intrinsic image decomposition algorithms.

306, TITLE: Planar Shape Detection at Structural Scales
http://openaccess.thecvf.com/content_cvpr_2018/html/Fang_Planar_Shape_Detection_CVPR_2018_paper.html
AUTHORS: Hao Fang, Florent Lafarge, Mathieu Desbrun
HIGHLIGHT: We present a framework to automatically extract a set of representations that capture the shape and structure of man-made objects at different key abstraction levels.

307, TITLE: Pix3D: Dataset and Methods for Single-Image 3D Shape Modeling
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Pix3D_Dataset_and_CVPR_2018_paper.html
AUTHORS: Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B. Tenenbaum, William T. Freeman
HIGHLIGHT: We study 3D shape modeling from a single image and make contributions to it in three aspects.

308, TITLE: Camera Pose Estimation With Unknown Principal Point
http://openaccess.thecvf.com/content_cvpr_2018/html/Larsson_Camera_Pose_Estimation_CVPR_2018_paper.html
AUTHORS: Viktor Larsson, Zuzana Kukelova, Yinqiang Zheng
HIGHLIGHT: In this paper, we develop the first exactly minimal solver for the case of unknown principal point and focal length by using four and a half point correspondences (P4.5Pfov).

- 309, TITLE: Inverse Composition Discriminative Optimization for Point Cloud Registration
http://openaccess.thecvf.com/content_cvpr_2018/html/Vongkulbhisal_Inverse_Composition_Discriminative_CVPR_2018_paper.html
AUTHORS: Jayakorn Vongkulbhisal, Beñat Irastorza Ugalde, Fernando De la Torre, João P. Costeira
HIGHLIGHT: To alleviate these issues, this paper proposes Inverse Composition Discriminative Optimization (ICDO), an extension of Discriminative Optimization (DO), which learns a sequence of update steps from synthetic training data that search the parameter space for an improved solution.
- 310, TITLE: SurfConv: Bridging 3D and 2D Convolution for RGBD Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Chu_SurfConv_Bridging_3D_CVPR_2018_paper.html
AUTHORS: Hang Chu, Wei-Chiu Ma, Kaustav Kundu, Raquel Urtasun, Sanja Fidler
HIGHLIGHT: Instead, we propose SurfConv, which “slides” compact 2D filters along the visible 3D surface.
- 311, TITLE: A Fast Resection-Intersection Method for the Known Rotation Problem
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_A_Fast_Resection-Intersection_CVPR_2018_paper.html
AUTHORS: Qianggong Zhang, Tat-Jun Chin, Huu Minh Le
HIGHLIGHT: In this paper, we devise a fast algorithm for the known rotation problem.
- 312, TITLE: 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Grabner_3D_Pose_Estimation_CVPR_2018_paper.html
AUTHORS: Alexander Grabner, Peter M. Roth, Vincent Lepetit
HIGHLIGHT: We propose a scalable, efficient and accurate approach to retrieve 3D models for objects in the wild.
- 313, TITLE: Structure From Recurrent Motion: From Rigidity to Recurrency
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Structure_From_Recurrent_CVPR_2018_paper.html
AUTHORS: Xiu Li, Hongdong Li, Hanbyul Joo, Yebin Liu, Yaser Sheikh
HIGHLIGHT: This paper proposes a new method for Non-rigidstructure-from-motion (NRSfM).
- 314, TITLE: Learning Patch Reconstructability for Accelerating Multi-View Stereo
http://openaccess.thecvf.com/content_cvpr_2018/html/Poms_Learning_Patch_Reconstructability_CVPR_2018_paper.html
AUTHORS: Alex Poms, Chenglei Wu, Shou-I Yu, Yaser Sheikh
HIGHLIGHT: We present an approach to accelerate multi-view stereo (MVS) by prioritizing computation on image patches that are likely to produce accurate 3D surface reconstructions.
- 315, TITLE: Progressively Complementarity-Aware Fusion Network for RGB-D Salient Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Progressively_Complementarity-Aware_Fusion_CVPR_2018_paper.html
AUTHORS: Hao Chen, Youfu Li
HIGHLIGHT: In this paper, we answer this question from two perspectives: (1) We argue that if the complementary part can be modelled more explicitly, the cross-modal complement is likely to be better captured.
- 316, TITLE: Pixels, Voxels, and Views: A Study of Shape Representations for Single View 3D Object Shape Prediction
http://openaccess.thecvf.com/content_cvpr_2018/html/Shin_Pixels_Voxels_and_CVPR_2018_paper.html
AUTHORS: Daeyun Shin, Charless C. Fowlkes, Derek Hoiem
HIGHLIGHT: The goal of this paper is to compare surface-based and volumetric 3D object shape representations, as well as viewer-centered and object-centered reference frames for single-view 3D shape prediction.
- 317, TITLE: Learning Dual Convolutional Neural Networks for Low-Level Vision
http://openaccess.thecvf.com/content_cvpr_2018/html/Pan_Learning_Dual_Convolutional_CVPR_2018_paper.html
AUTHORS: Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, Yu-Wing Tai, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a general dual convolutional neural network (DualCNN) for low-level vision problems, e.g., super-resolution, edge-preserving filtering, deraining and dehazing.
- 318, TITLE: Defocus Blur Detection via Multi-Stream Bottom-Top-Bottom Fully Convolutional Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Defocus_Blur_Detection_CVPR_2018_paper.html
AUTHORS: Wenda Zhao, Fan Zhao, Dong Wang, Huchuan Lu
HIGHLIGHT: To address these issues, we propose a multi-stream bottom-top-bottom fully convolutional network (BTBNet), which is the first attempt to develop an end-to-end deep network for DBD.
To promote further study and evaluation of the DBD models, we construct a new database of 500 challenging images and their pixel-wise defocus blur annotations.

- 319, TITLE: PiCANet: Learning Pixel-Wise Contextual Attention for Saliency Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_PiCANet_Learning_Pixel-Wise_CVPR_2018_paper.html
AUTHORS: Nian Liu, Junwei Han, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a novel pixel-wise contextual attention network, i.e., the PiCANet, to learn to selectively attend to informative context locations for each pixel.
- 320, TITLE: Curve Reconstruction via the Global Statistics of Natural Curves
http://openaccess.thecvf.com/content_cvpr_2018/html/Barnea_Curve_Reconstruction_via_CVPR_2018_paper.html
AUTHORS: Ehud Barnea, Ohad Ben-Shahar
HIGHLIGHT: Reconstructing the missing parts of a curve has been the subject of much computational research, with applications in image inpainting, object synthesis, etc.
- 321, TITLE: What Do Deep Networks Like to See?
http://openaccess.thecvf.com/content_cvpr_2018/html/Palacio_What_Do_Deep_CVPR_2018_paper.html
AUTHORS: Sebastian Palacio, Joachim Folz, Jörn Hees, Federico Raue, Damian Borth, Andreas Dengel
HIGHLIGHT: We propose a novel way to measure and understand convolutional neural networks by quantifying the amount of input signal they let in.
- 322, TITLE: "Zero-Shot" Super-Resolution Using Deep Internal Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Shocher_Zero-Shot_Super-Resolution_Using_CVPR_2018_paper.html
AUTHORS: Assaf Shocher, Nadav Cohen, Michal Irani
HIGHLIGHT: In this paper we introduce "Zero-Shot" SR, which exploits the power of Deep Learning, but does not rely on prior training.
- 323, TITLE: Detect Globally, Refine Locally: A Novel Approach to Saliency Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Detect_Globally_Refine_CVPR_2018_paper.html
AUTHORS: Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, Ali Borji
HIGHLIGHT: To address this problem, we propose a global Recurrent Localization Network (RLN) which exploits contextual information by the weighted response map in order to localize salient objects more accurately.
- 324, TITLE: Beyond the Pixel-Wise Loss for Topology-Aware Delineation
http://openaccess.thecvf.com/content_cvpr_2018/html/Mosinska_Beyond_the_Pixel-Wise_CVPR_2018_paper.html
AUTHORS: Agata Mosinska, Pablo Márquez-Neila, Mateusz Koziński, Pascal Fua
HIGHLIGHT: In this paper we claim that pixel-wise losses alone are unsuitable for this problem because of their inability to reflect the topological importance of prediction errors.
- 325, TITLE: KIPPI: KInetic Polygonal Partitioning of Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Bauchet_KIPPI_KInetic_Polygonal_CVPR_2018_paper.html
AUTHORS: Jean-Philippe Bauchet, Florent Lafarge
HIGHLIGHT: We propose a kinetic approach that brings more flexibility on polygon shape and size.
- 326, TITLE: Image Blind Denoising With Generative Adversarial Network Based Noise Modeling
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Image_Blind_Denoising_CVPR_2018_paper.html
AUTHORS: Jingwen Chen, Jiawei Chen, Hongyang Chao, Ming Yang
HIGHLIGHT: In this paper, we consider a typical image blind denoising problem, which is to remove unknown noise from noisy images.
- 327, TITLE: Multi-Scale Weighted Nuclear Norm Image Restoration
http://openaccess.thecvf.com/content_cvpr_2018/html/Yair_Multi-Scale_Weighted_Nuclear_CVPR_2018_paper.html
AUTHORS: Noam Yair, Tomer Michaeli
HIGHLIGHT: In this paper, we extend the WNNM method into a general image restoration algorithm, capable of handling arbitrary degradations (e.g. blur, missing pixels, etc.).
- 328, TITLE: MoNet: Moments Embedding Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Gou_MoNet_Moments_Embedding_CVPR_2018_paper.html
AUTHORS: Mengran Gou, Fei Xiong, Octavia Camps, Mario Sznajder
HIGHLIGHT: In this paper, we unify bilinear pooling and the global Gaussian embedding layers through the empirical moment matrix.

- 329, TITLE: Active Fixation Control to Predict Saccade Sequences
http://openaccess.thecvf.com/content_cvpr_2018/html/Wloka_Active_Fixation_Control_CVPR_2018_paper.html
AUTHORS: Calden Wloka, Iuliia Kotseruba, John K. Tsotsos
HIGHLIGHT: Towards addressing these shortcomings we present STAR-FC, a novel multi-saccade generator based on the integration of central high-level and object-based saliency and peripheral lower-level feature-based saliency.
- 330, TITLE: Densely Connected Pyramid Dehazing Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Densely_Connected_Pyramid_CVPR_2018_paper.html
AUTHORS: He Zhang, Vishal M. Patel
HIGHLIGHT: We propose a new end-to-end single image dehazing method, called Densely Connected Pyramid Dehazing Network (DCPDN), which can jointly learn the transmission map, atmospheric light and dehazing all together.
- 331, TITLE: Universal Denoising Networks : A Novel CNN Architecture for Image Denoising
http://openaccess.thecvf.com/content_cvpr_2018/html/Lefkimmiatis_Universal_Denoising_Networks_CVPR_2018_paper.html
AUTHORS: Stamatios Lefkimmiatis
HIGHLIGHT: Based on the proposed architecture, we introduce two different variants.
- 332, TITLE: Learning Convolutional Networks for Content-Weighted Image Compression
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Learning_Convolutional_Networks_CVPR_2018_paper.html
AUTHORS: Mu Li, Wangmeng Zuo, Shuhang Gu, Debin Zhao, David Zhang
HIGHLIGHT: In this paper, motivated by that the local information content is spatially variant in an image, we suggest that: (i) the bit rate of the different parts of the image is adapted to local content, and (ii) the content-aware bit rate is allocated under the guidance of a content-weighted importance map.
- 333, TITLE: Deep Video Super-Resolution Network Using Dynamic Upsampling Filters Without Explicit Motion Compensation
http://openaccess.thecvf.com/content_cvpr_2018/html/Jo_Deep_Video_Super-Resolution_CVPR_2018_paper.html
AUTHORS: Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, Seon Joo Kim
HIGHLIGHT: We introduce a fundamentally different framework for VSR in this paper.
- 334, TITLE: Erase or Fill? Deep Joint Recurrent Rain Removal and Reconstruction in Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Erase_or_Fill_CVPR_2018_paper.html
AUTHORS: Jiaying Liu, Wenhan Yang, Shuai Yang, Zongming Guo
HIGHLIGHT: In this paper, we address the problem of video rain removal by constructing deep recurrent convolutional networks.
- 335, TITLE: Flow Guided Recurrent Neural Encoder for Video Salient Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Flow_Guided_Recurrent_CVPR_2018_paper.html
AUTHORS: Guanbin Li, Yuan Xie, Tianhao Wei, Keze Wang, Liang Lin
HIGHLIGHT: In this paper, we present flow guided recurrent neural encoder(FGRNE), an accurate and end-to-end learning framework for video salient object detection.
- 336, TITLE: Gated Fusion Network for Single Image Dehazing
http://openaccess.thecvf.com/content_cvpr_2018/html/Ren_Gated_Fusion_Network_CVPR_2018_paper.html
AUTHORS: Wenqi Ren, Lin Ma, Jiawei Zhang, Jinshan Pan, Xiaochun Cao, Wei Liu, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose an efficient algorithm to directly restore a clear image from a hazy input.
- 337, TITLE: Learning a Single Convolutional Super-Resolution Network for Multiple Degradations
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Learning_a_Single_CVPR_2018_paper.html
AUTHORS: Kai Zhang, Wangmeng Zuo, Lei Zhang
HIGHLIGHT: To address these issues, we propose a general framework with dimensionality stretching strategy that enables a single convolutional super-resolution network to take two key factors of the SISR degradation process, i.e., blur kernel and noise level, as input.
- 338, TITLE: Non-Blind Deblurring: Handling Kernel Uncertainty With CNNs
http://openaccess.thecvf.com/content_cvpr_2018/html/Vasu_Non-Blind_Deblurring_Handling_CVPR_2018_paper.html
AUTHORS: Subeesh Vasu, Venkatesh Reddy Maligireddy, A. N. Rajagopalan
HIGHLIGHT: In this work, we present a convolutional neural network-based approach to handle kernel uncertainty in non-blind motion deblurring.

- 339, TITLE: Boundary Flow: A Siamese Network That Predicts Boundary Motion Without Training on Motion
http://openaccess.thecvf.com/content_cvpr_2018/html/Lei_Boundary_Flow_A_CVPR_2018_paper.html
AUTHORS: Peng Lei, Fuxin Li, Sinisa Todorovic
HIGHLIGHT: On boundary flow estimation, we present the first results on the Sintel training dataset.
- 340, TITLE: Learning to See in the Dark
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Learning_to_See_CVPR_2018_paper.html
AUTHORS: Chen Chen, Qifeng Chen, Jia Xu, Vladlen Koltun
HIGHLIGHT: To support the development of learning-based pipelines for low-light image processing, we introduce a dataset of raw short-exposure low-light images, with corresponding long-exposure reference images.
- 341, TITLE: BPGrad: Towards Global Optimality in Deep Learning via Branch and Pruning
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_BPGrad_Towards_Global_CVPR_2018_paper.html
AUTHORS: Ziming Zhang, Yuanwei Wu, Guanghui Wang
HIGHLIGHT: In this paper we propose a novel approximation algorithm, {em BPGrad}, towards optimizing deep models globally via branch and pruning.
- 342, TITLE: Perturbative Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Juefei-Xu_Perturbative_Neural_Networks_CVPR_2018_paper.html
AUTHORS: Felix Juefei-Xu, Vishnu Naresh Boddeti, Marios Savvides
HIGHLIGHT: In this paper, we seek to revisit the convolutional layer that has been the workhorse of state-of-the-art visual recognition models.
- 343, TITLE: Unsupervised Correlation Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Hoshen_Unsupervised_Correlation_Analysis_CVPR_2018_paper.html
AUTHORS: Yedid Hoshen, Lior Wolf
HIGHLIGHT: In this paper, we set to answer a fundamental cognitive question: are prior correspondences necessary for linking between different domains?
- 344, TITLE: A Biresolution Spectral Framework for Product Quantization
http://openaccess.thecvf.com/content_cvpr_2018/html/Mukherjee_A_Biresolution_Spectral_CVPR_2018_paper.html
AUTHORS: Lopamudra Mukherjee, Sathya N. Ravi, Jiming Peng, Vikas Singh
HIGHLIGHT: In this paper, we study the quantization problem in the setting where subspaces are orthogonal and show that this problem is intricately related to a specific type of spectral decomposition of the data.
- 345, TITLE: Domain Adaptive Faster R-CNN for Object Detection in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Domain_Adaptive_Faster_CVPR_2018_paper.html
AUTHORS: Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, Luc Van Gool
HIGHLIGHT: In this work, we aim to improve the cross-domain robustness of object detection.
- 346, TITLE: Low-Shot Learning With Large-Scale Diffusion
http://openaccess.thecvf.com/content_cvpr_2018/html/Douze_Low-Shot_Learning_With_CVPR_2018_paper.html
AUTHORS: Matthijs Douze, Arthur Szlam, Bharath Hariharan, Hervé Jégou
HIGHLIGHT: This paper considers the problem of inferring image labels from images when only a few annotated examples are available at training time.
- 347, TITLE: Joint Pose and Expression Modeling for Facial Expression Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Joint_Pose_and_CVPR_2018_paper.html
AUTHORS: Feifei Zhang, Tianzhu Zhang, Qirong Mao, Changsheng Xu
HIGHLIGHT: Different from existing methods, in this paper, we propose an end-to-end deep learning model by exploiting different poses and expressions jointly for simultaneous facial image synthesis and pose-invariant facial expression recognition.
- 348, TITLE: Lightweight Probabilistic Deep Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Gast_Lightweight_Probabilistic_Deep_CVPR_2018_paper.html
AUTHORS: Jochen Gast, Stefan Roth
HIGHLIGHT: Lightweight Probabilistic Deep Networks
- 349, TITLE: Adversarially Learned One-Class Classifier for Novelty Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Sabokrou_Adversarially_Learned_One-Class_CVPR_2018_paper.html
AUTHORS: Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, Ehsan Adeli

HIGHLIGHT: In this paper, inspired by the success of generative adversarial networks for training deep models in unsupervised and semi-supervised settings, we propose an end-to-end architecture for one-class classification.

350, **TITLE:** Defense Against Universal Adversarial Perturbations
http://openaccess.thecvf.com/content_cvpr_2018/html/Akhtar_Defense_Against_Universal_CVPR_2018_paper.html
AUTHORS: Naveed Akhtar, Jian Liu, Ajmal Mian
HIGHLIGHT: We present the first dedicated framework to effectively defend the networks against such perturbations.

351, **TITLE:** Disentangling Factors of Variation by Mixing Them
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Disentangling_Factors_of_CVPR_2018_paper.html
AUTHORS: Qiyang Hu, Attila Szabó, Tiziano Portenier, Paolo Favaro, Matthias Zwicker
HIGHLIGHT: We propose an approach to learn image representations that consist of disentangled factors of variation without exploiting any manual labeling or data domain knowledge.

352, **TITLE:** Deformable GANs for Pose-Based Human Image Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Siarohin_Deformable_GANs_for_CVPR_2018_paper.html
AUTHORS: Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, Nicu Sebe
HIGHLIGHT: In this paper we address the problem of generating person images conditioned on a given pose.

353, **TITLE:** Hierarchical Recurrent Attention Networks for Structured Online Maps
http://openaccess.thecvf.com/content_cvpr_2018/html/Homayounfar_Hierarchical_Recurrent_Attention_CVPR_2018_paper.html
AUTHORS: Namdar Homayounfar, Wei-Chiu Ma, Shrinidhi Kowshika Lakshmikanth, Raquel Urtasun
HIGHLIGHT: In this paper, we tackle the problem of online road network extraction from sparse 3D point clouds.

354, **TITLE:** Sliced Wasserstein Distance for Learning Gaussian Mixture Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Kolouri_Sliced_Wasserstein_Distance_CVPR_2018_paper.html
AUTHORS: Soheil Kolouri, Gustavo K. Rohde, Heiko Hoffmann
HIGHLIGHT: Inspired by the relationship between the negative log-likelihood function and the Kullback-Leibler (KL) divergence, we propose an alternative formulation for estimating the GMM parameters using the sliced Wasserstein distance, which gives rise to a new algorithm.

355, **TITLE:** Aligning Infinite-Dimensional Covariance Matrices in Reproducing Kernel Hilbert Spaces for Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Aligning_Infinite-Dimensional_Covariance_CVPR_2018_paper.html
AUTHORS: Zhen Zhang, Mianzhi Wang, Yan Huang, Arye Nehorai
HIGHLIGHT: In this paper, we provide two alignment approaches, for both of which we obtain closed-form expressions via kernel matrices.

356, **TITLE:** CLEAR: Cumulative LEARning for One-Shot One-Class Image Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Kozeraowski_CLEAR_Cumulative_LEARning_CVPR_2018_paper.html
AUTHORS: Jędrzej Kozeraowski, Matthew Turk
HIGHLIGHT: This work addresses the novel problem of one-shot one-class classification.

357, **TITLE:** Local and Global Optimization Techniques in Graph-Based Clustering
http://openaccess.thecvf.com/content_cvpr_2018/html/ikami_Local_and_Global_CVPR_2018_paper.html
AUTHORS: Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa
HIGHLIGHT: We propose a local optimization method, which is widely applicable to graph-based clustering cost functions.

358, **TITLE:** Multi-Task Learning by Maximizing Statistical Dependence
http://openaccess.thecvf.com/content_cvpr_2018/html/Mejjati_Multi-Task_Learning_by_CVPR_2018_paper.html
AUTHORS: Youssef A. Mejjati, Darren Cosker, Kwang In Kim
HIGHLIGHT: We present a new multi-task learning (MTL) approach that can be applied to multiple heterogeneous task estimators.

359, **TITLE:** Robust Classification With Convolutional Prototype Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Robust_Classification_With_CVPR_2018_paper.html
AUTHORS: Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, Cheng-Lin Liu
HIGHLIGHT: In this paper, we argue that the lack of robustness for CNN is caused by the softmax layer, which is a totally discriminative model and based on the assumption of closed world (i.e., with a fixed number of categories).

- 360, TITLE: Generative Modeling Using the Sliced Wasserstein Distance
http://openaccess.thecvf.com/content_cvpr_2018/html/Deshpande_Generative_Modeling_Using_CVPR_2018_paper.html
AUTHORS: Ishan Deshpande, Ziyu Zhang, Alexander G. Schwing
HIGHLIGHT: By augmenting this approach with a discriminator we improve its accuracy.
- 361, TITLE: Learning Time/Memory-Efficient Deep Architectures With Budgeted Super Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Veniat_Learning_TimeMemory-Efficient_Deep_CVPR_2018_paper.html
AUTHORS: Tom Véniat, Ludovic Denoyer
HIGHLIGHT: We propose to focus on the problem of discovering neural network architectures efficient in terms of both prediction quality and cost.
We present a set of experiments on computer vision problems and analyze the ability of our technique to deal with three different costs: the computation cost, the memory consumption cost and a distributed computation cost.
- 362, TITLE: Cross-View Image Synthesis Using Conditional GANs
http://openaccess.thecvf.com/content_cvpr_2018/html/Regmi_Cross-View_Image_Synthesis_CVPR_2018_paper.html
AUTHORS: Krishna Regmi, Ali Borji
HIGHLIGHT: In this paper, we attempt to solve the novel problem of cross-view image synthesis, aerial to street-view and vice versa, using conditional generative adversarial networks (cGAN).
- 363, TITLE: Sparse, Smart Contours to Represent and Edit Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Dekel_Sparse_Smart_Contours_CVPR_2018_paper.html
AUTHORS: Tali Dekel, Chuang Gan, Dilip Krishnan, Ce Liu, William T. Freeman
HIGHLIGHT: We study the problem of reconstructing an image from information stored at contour locations.
- 364, TITLE: Anticipating Traffic Accidents With Adaptive Loss and Large-Scale Incident DB
http://openaccess.thecvf.com/content_cvpr_2018/html/Suzuki_Anticipating_Traffic_Accidents_CVPR_2018_paper.html
AUTHORS: Tomoyuki Suzuki, Hirokatsu Kataoka, Yoshimitsu Aoki, Yutaka Satoh
HIGHLIGHT: In this paper, we propose a novel approach for traffic accident anticipation through (i) Adaptive Loss for Early Anticipation (AdaLEA) and (ii) a large-scale self-annotated incident database.
- 365, TITLE: A Minimalist Approach to Type-Agnostic Detection of Quadrics in Point Clouds
http://openaccess.thecvf.com/content_cvpr_2018/html/Birdal_A_Minimalist_Approach_CVPR_2018_paper.html
AUTHORS: Tolga Birdal, Benjamin Busam, Nassir Navab, Slobodan Ilic, Peter Sturm
HIGHLIGHT: This paper proposes a segmentation-free, automatic and efficient procedure to detect general geometric quadric forms in point clouds, where clutter and occlusions are inevitable.
- 366, TITLE: Facelet-Bank for Fast Portrait Manipulation
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Facelet-Bank_for_Fast_CVPR_2018_paper.html
AUTHORS: Ying-Cong Chen, Huaijia Lin, Michelle Shu, Ruiyu Li, Xin Tao, Xiaoyong Shen, Yangang Ye, Jiaya Jia
HIGHLIGHT: In this paper, we propose a model to achieve this goal based on an end-to-end convolutional neural network that supports fast inference, edit-effect control, and quick partial-model update.
- 367, TITLE: Visual to Sound: Generating Natural Sound for Videos in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Visual_to_Sound_CVPR_2018_paper.html
AUTHORS: Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, Tamara L. Berg
HIGHLIGHT: In this paper, we pose the task of generating sound given visual input.
- 368, TITLE: 3D-RCNN: Instance-Level 3D Object Reconstruction via Render-and-Compare
http://openaccess.thecvf.com/content_cvpr_2018/html/Kundu_3D-RCNN_Instance-Level_3D_CVPR_2018_paper.html
AUTHORS: Abhijit Kundu, Yin Li, James M. Rehg
HIGHLIGHT: We present a fast inverse-graphics framework for instance-level 3D scene understanding.
- 369, TITLE: Fast and Furious: Real Time End-to-End 3D Detection, Tracking and Motion Forecasting With a Single Convolutional Net
http://openaccess.thecvf.com/content_cvpr_2018/html/Luo_Fast_and_Furious_CVPR_2018_paper.html
AUTHORS: Wenjie Luo, Bin Yang, Raquel Urtasun
HIGHLIGHT: In this paper we propose a novel deep neural network that is able to jointly reason about 3D detection, tracking and motion forecasting given data captured by a 3D sensor.

- 370, TITLE: An Analysis of Scale Invariance in Object Detection - SNIP
http://openaccess.thecvf.com/content_cvpr_2018/html/Singh_An_Analysis_of_CVPR_2018_paper.html
AUTHORS: Bharat Singh, Larry S. Davis
HIGHLIGHT: Based on this analysis, we propose to train and test detectors on the same scales of an image-pyramid.
- 371, TITLE: Relation Networks for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Relation_Networks_for_CVPR_2018_paper.html
AUTHORS: Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, Yichen Wei
HIGHLIGHT: This work proposes an object relation module.
- 372, TITLE: Zero-Shot Sketch-Image Hashing
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Zero-Shot_Sketch-Image_Hashing_CVPR_2018_paper.html
AUTHORS: Yuming Shen, Li Liu, Fumin Shen, Ling Shao
HIGHLIGHT: In this paper, the above problem is briefed as a novel but realistic zero-shot SBIR hashing task.
- 373, TITLE: VizWiz Grand Challenge: Answering Visual Questions From Blind People
http://openaccess.thecvf.com/content_cvpr_2018/html/Gurari_VizWiz_Grand_Challenge_CVPR_2018_paper.html
AUTHORS: Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, Jeffrey P. Bigham
HIGHLIGHT: We introduce this dataset to encourage a larger community to develop more generalized algorithms that can assist blind people.
- 374, TITLE: Divide and Grow: Capturing Huge Diversity in Crowd Images With Incrementally Growing CNN
http://openaccess.thecvf.com/content_cvpr_2018/html/Sam_Divide_and_Grow_CVPR_2018_paper.html
AUTHORS: Deepak Babu Sam, Neeraj N. Sajjan, R. Venkatesh Babu, Mukundhan Srinivasan
HIGHLIGHT: We tackle this problem with a growing CNN which can progressively increase its capacity to account for the wide variability seen in crowd scenes.
- 375, TITLE: Structured Set Matching Networks for One-Shot Part Labeling
http://openaccess.thecvf.com/content_cvpr_2018/html/Choi_Structured_Set_Matching_CVPR_2018_paper.html
AUTHORS: Jonghyun Choi, Jayant Krishnamurthy, Aniruddha Kembhavi, Ali Farhadi
HIGHLIGHT: For this set-to-set matching problem, we introduce the Structured Set Matching Network (SSMN), a structured prediction model that incorporates convolutional neural networks.
- 376, TITLE: Self-Supervised Learning of Geometrically Stable Features Through Probabilistic Introspection
http://openaccess.thecvf.com/content_cvpr_2018/html/Novotny_Self-Supervised_Learning_of_CVPR_2018_paper.html
AUTHORS: David Novotny, Samuel Albanie, Diane Larlus, Andrea Vedaldi
HIGHLIGHT: Our approach learns dense distinctive visual descriptors from an unlabeled dataset of images using synthetic image transformations.
- 377, TITLE: Link and Code: Fast Indexing With Graphs and Compact Regression Codes
http://openaccess.thecvf.com/content_cvpr_2018/html/Douze_Link_and_Code_CVPR_2018_paper.html
AUTHORS: Matthijs Douze, Alexandre Sablayrolles, Hervé Jégou
HIGHLIGHT: In this paper, we revisit these approaches by considering, additionally, the memory constraint required to index billions of images on a single server.
- 378, TITLE: Textbook Question Answering Under Instructor Guidance With Memory Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Textbook_Question_Answering_CVPR_2018_paper.html
AUTHORS: Juzheng Li, Hang Su, Jun Zhu, Siyu Wang, Bo Zhang
HIGHLIGHT: To address this issue, we propose a novel approach of Instructor Guidance with Memory Networks (IGMN) which conducts the TQA task by finding contradictions between the candidate answers and their corresponding context.
- 379, TITLE: Unsupervised Deep Generative Adversarial Hashing Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Dizaji_Unsupervised_Deep_Generative_CVPR_2018_paper.html
AUTHORS: Kamran Ghasedi Dizaji, Feng Zheng, Najmeh Sadoughi, Yanhua Yang, Cheng Deng, Heng Huang
HIGHLIGHT: In this paper, we propose a deep unsupervised hashing function, called HashGAN, which outperforms unsupervised hashing models with significant margins without any supervised pretraining.
- 380, TITLE: Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments

http://openaccess.thecvf.com/content_cvpr_2018/html/Anderson_Vision-and-Language_Navigation_Interpreting_CVPR_2018_paper.html

AUTHORS: Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, Anton van den Hengel

HIGHLIGHT: To enable and encourage the application of vision and language methods to the problem of interpreting visually-grounded navigation instructions, we present the Matterport3D Simulator -- a large-scale reinforcement learning environment based on real imagery.

Using this simulator, which can in future support a range of embodied vision and language tasks, we provide the first benchmark dataset for visually-grounded natural language navigation in real buildings -- the Room-to-Room (R2R) dataset.

381, TITLE: DenseASPP for Semantic Segmentation in Street Scenes

http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_DenseASPP_for_Semantic_CVPR_2018_paper.html

AUTHORS: Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, Kuiyuan Yang

HIGHLIGHT: To this end, we propose Densely connected Atrous Spatial Pyramid Pooling (DenseASPP), which connects a set of atrous convolutional layers in a dense way, such that it generates multi-scale features that not only cover a larger scale range, but also cover that scale range densely, without significantly increasing the model size.

382, TITLE: Efficient Optimization for Rank-Based Loss Functions

http://openaccess.thecvf.com/content_cvpr_2018/html/Mohapatra_Efficient_Optimization_for_CVPR_2018_paper.html

AUTHORS: Pritish Mohapatra, Michal Rolínek, C.V. Jawahar, Vladimir Kolmogorov, M. Pawan Kumar

HIGHLIGHT: To alleviate this deficiency, we present a novel quicksort flavored algorithm for a large class of non-decomposable loss functions.

383, TITLE: Wasserstein Introspective Neural Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_Wasserstein_Introspective_Neural_CVPR_2018_paper.html

AUTHORS: Kwonjoon Lee, Weijian Xu, Fan Fan, Zhuowen Tu

HIGHLIGHT: We present Wasserstein introspective neural networks (WINN) that are both a generator and a discriminator within a single model.

384, TITLE: Taskonomy: Disentangling Task Transfer Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Zamir_Taskonomy_Disentangling_Task_CVPR_2018_paper.html

AUTHORS: Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, Silvio Savarese

HIGHLIGHT: This paper proposes a fully computational approach for finding the structure of the space of visual tasks.

385, TITLE: Maximum Classifier Discrepancy for Unsupervised Domain Adaptation

http://openaccess.thecvf.com/content_cvpr_2018/html/Saito_Maximum_Classifier_Discrepancy_CVPR_2018_paper.html

AUTHORS: Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, Tatsuya Harada

HIGHLIGHT: In this work, we present a method for unsupervised domain adaptation.

386, TITLE: Unsupervised Feature Learning via Non-Parametric Instance Discrimination

http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Unsupervised_Feature_Learning_CVPR_2018_paper.html

AUTHORS: Zhirong Wu, Yuanjun Xiong, Stella X. Yu, Dahua Lin

HIGHLIGHT: We formulate this intuition as a non-parametric classification problem at the instance-level, and use noise-contrastive estimation to tackle the computational challenges imposed by the large number of instance classes.

387, TITLE: Multi-Task Adversarial Network for Disentangled Feature Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Multi-Task_Adversarial_Network_CVPR_2018_paper.html

AUTHORS: Yang Liu, Zhaowen Wang, Hailin Jin, Ian Wassell

HIGHLIGHT: We present a novel multi-task adversarial network based on an encoder-discriminator-generator architecture.

388, TITLE: Learning From Synthetic Data: Addressing Domain Shift for Semantic Segmentation

http://openaccess.thecvf.com/content_cvpr_2018/html/Sankaranarayanan_Learning_From_Synthetic_CVPR_2018_paper.html

AUTHORS: Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, Rama Chellappa

HIGHLIGHT: In this work, we focus on adapting the representations learned by segmentation networks across synthetic and real domains.

389, TITLE: Empirical Study of the Topology and Geometry of Deep Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Fawzi_Empirical_Study_of_CVPR_2018_paper.html

AUTHORS: Alhussein Fawzi, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard, Stefano Soatto

HIGHLIGHT: The goal of this paper is to analyze the geometric properties of deep neural network image classifiers in the input space.

- 390, TITLE: Boosting Domain Adaptation by Discovering Latent Domains
http://openaccess.thecvf.com/content_cvpr_2018/html/Mancini_Boosting_Domain_Adaptation_CVPR_2018_paper.html
AUTHORS: Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, Elisa Ricci
HIGHLIGHT: This paper introduces a novel Convolutional Neural Network (CNN) architecture which (i) automatically discovers latent domains in visual datasets and (ii) exploits this information to learn robust target classifiers.
- 391, TITLE: Shape From Shading Through Shape Evolution
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Shape_From_Shading_CVPR_2018_paper.html
AUTHORS: Dawei Yang, Jia Deng
HIGHLIGHT: In this paper, we address the shape-from-shading problem by training deep networks with synthetic images.
- 392, TITLE: Weakly Supervised Instance Segmentation Using Class Peak Response
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Weakly_Supervised_Instance_CVPR_2018_paper.html
AUTHORS: Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, Jianbin Jiao
HIGHLIGHT: In this paper, we tackle this challenging problem by exploiting class peak responses to enable a classification network for instance mask extraction.
- 393, TITLE: Collaborative and Adversarial Network for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Collaborative_and_Adversarial_CVPR_2018_paper.html
AUTHORS: Weichen Zhang, Wanli Ouyang, Wen Li, Dong Xu
HIGHLIGHT: In this paper, we propose a new unsupervised domain adaptation approach called Collaborative and Adversarial Network (CAN) through domain-collaborative and domain-adversarial training of neural networks.
- 394, TITLE: Environment Upgrade Reinforcement Learning for Non-Differentiable Multi-Stage Pipelines
http://openaccess.thecvf.com/content_cvpr_2018/html/Xie_Environment_Upgrade_Reinforcement_CVPR_2018_paper.html
AUTHORS: Shuqin Xie, Zitian Chen, Chao Xu, Cewu Lu
HIGHLIGHT: In this paper, we propose a novel environment upgrade reinforcement learning framework to solve the feedback and joint optimization problems.
- 395, TITLE: Teaching Categories to Human Learners With Visual Explanations
http://openaccess.thecvf.com/content_cvpr_2018/html/Aodha_Teaching_Categories_to_CVPR_2018_paper.html
AUTHORS: Oisín Mac Aodha, Shihan Su, Yuxin Chen, Pietro Perona, Yisong Yue
HIGHLIGHT: To address these existing limitations, we propose a teaching framework that provides interpretable explanations as feedback and models how the learner incorporates this additional information.
- 396, TITLE: Density Adaptive Point Set Registration
http://openaccess.thecvf.com/content_cvpr_2018/html/Lawin_Density_Adaptive_Point_CVPR_2018_paper.html
AUTHORS: Felix Järemo Lawin, Martin Danelljan, Fahad Shahbaz Khan, Per-Erik Forssén, Michael Felsberg
HIGHLIGHT: Probabilistic methods for point set registration have demonstrated competitive results in recent years.
- 397, TITLE: Left-Right Comparative Recurrent Model for Stereo Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Jie_Left-Right_Comparative_Recurrent_CVPR_2018_paper.html
AUTHORS: Zequn Jie, Pengfei Wang, Yonggen Ling, Bo Zhao, Yunchao Wei, Jiashi Feng, Wei Liu
HIGHLIGHT: This paper proposes a novel left-right comparative recurrent model to perform left-right consistency checking jointly with disparity estimation.
- 398, TITLE: Im2Pano3D: Extrapolating 360° Structure and Semantics Beyond the Field of View
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_Im2Pano3D_Extrapolating_360deg_CVPR_2018_paper.html
AUTHORS: Shuran Song, Andy Zeng, Angel X. Chang, Manolis Savva, Silvio Savarese, Thomas Funkhouser
HIGHLIGHT: We present Im2Pano3D, a convolutional neural network that generates a dense prediction of 3D structure and a probability distribution of semantic labels for a full 360 panoramic view of an indoor scene when given only a partial observation ($\approx 50\%$) in the form of an RGB-D image.
- 399, TITLE: Polarimetric Dense Monocular SLAM
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Polarimetric_Dense_Monocular_CVPR_2018_paper.html
AUTHORS: Luwei Yang, Feitong Tan, Ao Li, Zhaopeng Cui, Yasutaka Furukawa, Ping Tan
HIGHLIGHT: This paper presents a novel polarimetric dense monocular SLAM (PDMS) algorithm based on a polarization camera.

- 400, TITLE: A Unifying Contrast Maximization Framework for Event Cameras, With Applications to Motion, Depth, and Optical Flow Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Gallego_A_Unifying_Contrast_CVPR_2018_paper.html
AUTHORS: Guillermo Gallego, Henri Rebecq, Davide Scaramuzza
HIGHLIGHT: We present a unifying framework to solve several computer vision problems with event cameras: motion, depth and optical flow estimation.
- 401, TITLE: Modeling Facial Geometry Using Compositional VAEs
http://openaccess.thecvf.com/content_cvpr_2018/html/Bagautdinov_Modeling_Facial_Geometry_CVPR_2018_paper.html
AUTHORS: Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, Yaser Sheikh
HIGHLIGHT: We propose a method for learning non-linear face geometry representations using deep generative models.
- 402, TITLE: Tangent Convolutions for Dense Prediction in 3D
http://openaccess.thecvf.com/content_cvpr_2018/html/Tatarchenko_Tangent_Convolutions_for_CVPR_2018_paper.html
AUTHORS: Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, Qian-Yi Zhou
HIGHLIGHT: We present an approach to semantic scene analysis using deep convolutional networks.
- 403, TITLE: RayNet: Learning Volumetric 3D Reconstruction With Ray Potentials
http://openaccess.thecvf.com/content_cvpr_2018/html/Paschalidou_RayNet_Learning_Volumetric_CVPR_2018_paper.html
AUTHORS: Despoina Paschalidou, Osman Ulusoy, Carolin Schmitt, Luc Van Gool, Andreas Geiger
HIGHLIGHT: In this paper, we consider the problem of reconstructing a dense 3D model using images captured from different views.
- 404, TITLE: Neural 3D Mesh Renderer
http://openaccess.thecvf.com/content_cvpr_2018/html/Kato_Neural_3D_Mesh_CVPR_2018_paper.html
AUTHORS: Hiroharu Kato, Yoshitaka Ushiku, Tatsuya Harada
HIGHLIGHT: Therefore, in this work, we propose an approximate gradient for rasterization that enables the integration of rendering into neural networks.
- 405, TITLE: Structured Attention Guided Convolutional Neural Fields for Monocular Depth Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Structured_Attention_Guided_CVPR_2018_paper.html
AUTHORS: Dan Xu, Wei Wang, Hao Tang, Hong Liu, Nicu Sebe, Elisa Ricci
HIGHLIGHT: Following this line of research, in this paper we introduce a novel approach for monocular depth estimation.
- 406, TITLE: Automatic 3D Indoor Scene Modeling From Single Panorama
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Automatic_3D_Indoor_CVPR_2018_paper.html
AUTHORS: Yang Yang, Shi Jin, Ruiyang Liu, Sing Bing Kang, Jingyi Yu
HIGHLIGHT: We describe a system that automatically extracts 3D geometry of an indoor scene from a single 2D panorama.
- 407, TITLE: Extreme 3D Face Reconstruction: Seeing Through Occlusions
http://openaccess.thecvf.com/content_cvpr_2018/html/Tran_Extreme_3D_Face_CVPR_2018_paper.html
AUTHORS: Anh Tu?n Tr?n, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, Gérard Medioni
HIGHLIGHT: Motivated by the concept of bump mapping, we propose a layered approach which decouples estimation of a global shape from its mid-level details (e.g., wrinkles).
- 408, TITLE: Beyond Grobner Bases: Basis Selection for Minimal Solvers
http://openaccess.thecvf.com/content_cvpr_2018/html/Larsson_Beyond_Grobner_Bases_CVPR_2018_paper.html
AUTHORS: Viktor Larsson, Magnus Oskarsson, Kalle Astrom, Alge Wallis, Zuzana Kukelova, Tomas Pajdla
HIGHLIGHT: In this paper we show how we can make polynomial solvers based on the action matrix method faster, by careful selection of the monomial bases.
- 409, TITLE: Lions and Tigers and Bears: Capturing Non-Rigid, 3D, Articulated Shape From Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Zuffi_Lions_and_Tigers_CVPR_2018_paper.html
AUTHORS: Silvia Zuffi, Angjoo Kanazawa, Michael J. Black
HIGHLIGHT: Consequently, we propose a method to capture the detailed 3D shape of animals from images alone.
- 410, TITLE: Deep Cocktail Network: Multi-Source Unsupervised Domain Adaptation With Category Shift
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Deep_Cocktail_Network_CVPR_2018_paper.html
AUTHORS: Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, Liang Lin

HIGHLIGHT: In this paper, we propose a deep cocktail network (DCTN), to battle the domain and category shifts among multiple sources.

411, **TITLE:** DOTA: A Large-Scale Dataset for Object Detection in Aerial Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Xia_DOTA_A_Large-Scale_CVPR_2018_paper.html
AUTHORS: Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, Liangpei Zhang
HIGHLIGHT: To advance object detection research in Earth Vision, also known as Earth Observation and Remote Sensing, we introduce a large-scale Dataset for Object deTection in Aerial images (DOTA).

412, **TITLE:** Finding Beans in Burgers: Deep Semantic-Visual Embedding With Localization
http://openaccess.thecvf.com/content_cvpr_2018/html/Engilberge_Finding_Beans_in_CVPR_2018_paper.html
AUTHORS: Martin Engilberge, Louis Chevallier, Patrick Pérez, Matthieu Cord
HIGHLIGHT: Several works have proposed to learn a two-path neural network that maps images and texts, respectively, to a same shared Euclidean space where geometry captures useful semantic relationships.

413, **TITLE:** Feature Super-Resolution: Make Machine See More Clearly
http://openaccess.thecvf.com/content_cvpr_2018/html/Tan_Feature_Super-Resolution_Make_CVPR_2018_paper.html
AUTHORS: Weimin Tan, Bo Yan, Bahetiyaer Bare
HIGHLIGHT: In this paper, different from image super-resolution (ISR), we propose a novel super-resolution technique called feature super-resolution (FSR), which aims at enhancing the discriminatory power of small size image in order to provide high recognition precision for machine.

414, **TITLE:** ClusterNet: Detecting Small Objects in Large Scenes by Exploiting Spatio-Temporal Information
http://openaccess.thecvf.com/content_cvpr_2018/html/LaLonde_ClusterNet_Detecting_Small_CVPR_2018_paper.html
AUTHORS: Rodney LaLonde, Dong Zhang, Mubarak Shah
HIGHLIGHT: In this work, we experimentally verify the failure of appearance-based classifiers in WAMI, such as Faster R-CNN and a heatmap-based fully convolutional neural network (CNN), and propose a novel two-stage spatio-temporal CNN which effectively and efficiently combines both appearance and motion information to significantly surpass the state-of-the-art in WAMI object detection.

415, **TITLE:** MaskLab: Instance Segmentation by Refining Object Detection With Semantic and Direction Features
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_MaskLab_Instance_Segmentation_CVPR_2018_paper.html
AUTHORS: Liang-Chieh Chen, Alexander Hermans, George Papandreou, Florian Schroff, Peng Wang, Hartwig Adam
HIGHLIGHT: In this work, we tackle the problem of instance segmentation, the task of simultaneously solving object detection and semantic segmentation.

416, **TITLE:** Hashing as Tie-Aware Learning to Rank
http://openaccess.thecvf.com/content_cvpr_2018/html/He_Hashing_as_Tie-Aware_CVPR_2018_paper.html
AUTHORS: Kun He, Fatih Cakir, Sarah Adel Bargal, Stan Sclaroff
HIGHLIGHT: In this paper, we develop learning to rank formulations for hashing, aimed at directly optimizing ranking-based evaluation metrics such as Average Precision (AP) and Normalized Discounted Cumulative Gain (NDCG).

417, **TITLE:** Classification-Driven Dynamic Image Enhancement
http://openaccess.thecvf.com/content_cvpr_2018/html/Sharma_Classification-Driven_Dynamic_Image_CVPR_2018_paper.html
AUTHORS: Vivek Sharma, Ali Diba, Davy Neven, Michael S. Brown, Luc Van Gool, Rainer Stiefelhagen
HIGHLIGHT: In this paper, we are interested in learning CNNs that can emulate image enhancement and restoration, but with the overall goal to improve image classification and not necessarily human perception.

418, **TITLE:** Knowledge Aided Consistency for Weakly Supervised Phrase Grounding
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Knowledge_Aided_Consistency_CVPR_2018_paper.html
AUTHORS: Kan Chen, Jiyang Gao, Ram Nevatia
HIGHLIGHT: In this paper, we explore the consistency contained in both visual and language modalities, and leverage complementary external knowledge to facilitate weakly supervised grounding.

419, **TITLE:** Who Let the Dogs Out? Modeling Dog Behavior From Visual Data
http://openaccess.thecvf.com/content_cvpr_2018/html/Ehsani_Who_Let_the_CVPR_2018_paper.html
AUTHORS: Kiana Ehsani, Hessam Bagherinezhad, Joseph Redmon, Roozbeh Mottaghi, Ali Farhadi
HIGHLIGHT: We introduce the task of directly modeling a visually intelligent agent.

- 420, TITLE: Pseudo Mask Augmented Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Pseudo_Mask_Augmented_CVPR_2018_paper.html
AUTHORS: Xiangyun Zhao, Shuang Liang, Yichen Wei
HIGHLIGHT: In this work, we present a novel and effective framework to facilitate object detection with the instance-level segmentation information that is only supervised by bounding box annotation.
- 421, TITLE: Dual Skipping Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_Dual_Skipping_Networks_CVPR_2018_paper.html
AUTHORS: Changmao Cheng, Yanwei Fu, Yu-Gang Jiang, Wei Liu, Wenlian Lu, Jianfeng Feng, Xiangyang Xue
HIGHLIGHT: Inspired by the recent neuroscience studies on the left-right asymmetry of the human brain in processing low and high spatial frequency information, this paper introduces a dual skipping network which carries out coarse-to-fine object categorization.
- 422, TITLE: Memory Matching Networks for One-Shot Image Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Cai_Memory_Matching_Networks_CVPR_2018_paper.html
AUTHORS: Qi Cai, Yingwei Pan, Ting Yao, Chenggang Yan, Tao Mei
HIGHLIGHT: In this paper, we introduce the new ideas of augmenting Convolutional Neural Networks (CNNs) with Memory and learning to learn the network parameters for the unlabelled images on the fly in one-shot learning.
- 423, TITLE: IQA: Visual Question Answering in Interactive Environments
http://openaccess.thecvf.com/content_cvpr_2018/html/Gordon_IQA_Visual_Question_CVPR_2018_paper.html
AUTHORS: Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, Ali Farhadi
HIGHLIGHT: We introduce Interactive Question Answering (IQA), the task of answering questions that require an autonomous agent to interact with a dynamic visual environment.
- 424, TITLE: Pose Transferrable Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Pose_Transferrable_Person_CVPR_2018_paper.html
AUTHORS: Jinxian Liu, Bingbing Ni, Yichao Yan, Peng Zhou, Shuo Cheng, Jianguo Hu
HIGHLIGHT: To address this issue, we propose a pose-transferrable person ReID framework which utilizes pose-transferred sample augmentations (i.e., with ID supervision) to enhance ReID model training.
- 425, TITLE: Large Scale Fine-Grained Categorization and Domain-Specific Transfer Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Cui_Large_Scale_Fine-Grained_CVPR_2018_paper.html
AUTHORS: Yin Cui, Yang Song, Chen Sun, Andrew Howard, Serge Belongie
HIGHLIGHT: We propose a measure to estimate domain similarity via Earth Mover's Distance and demonstrate that transfer learning benefits from pre-training on a source domain that is similar to the target domain by this measure.
- 426, TITLE: Data Distillation: Towards Omni-Supervised Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Radosavovic_Data_Distillation_Towards_CVPR_2018_paper.html
AUTHORS: Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, Kaiming He
HIGHLIGHT: To exploit the omni-supervised setting, we propose data distillation, a method that ensembles predictions from multiple transformations of unlabeled data, using a single model, to automatically generate new training annotations.
- 427, TITLE: Object Referring in Videos With Language and Human Gaze
http://openaccess.thecvf.com/content_cvpr_2018/html/Vasudevan_Object_Referring_in_CVPR_2018_paper.html
AUTHORS: Arun Balajee Vasudevan, Dengxin Dai, Luc Van Gool
HIGHLIGHT: We investigate the problem of object referring (OR) i.e. to localize a target object in a visual scene coming with a language description.
To that end, we present a new video dataset for OR, with 30,000 objects over 5,000 stereo video sequences annotated for their descriptions and gaze.
- 428, TITLE: Feature Selective Networks for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhai_Feature_Selective_Networks_CVPR_2018_paper.html
AUTHORS: Yao Zhai, Jingjing Fu, Yan Lu, Houqiang Li
HIGHLIGHT: We present feature selective networks to reform the feature representations of RoIs by exploiting their disparities among sub-regions and aspect ratios.
- 429, TITLE: Learning a Discriminative Filter Bank Within a CNN for Fine-Grained Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Learning_a_Discriminative_CVPR_2018_paper.html
AUTHORS: Yaming Wang, Vlad I. Morariu, Larry S. Davis

HIGHLIGHT: Compared to earlier multistage frameworks using CNN features, recent end-to-end deep approaches for fine-grained recognition essentially enhance the mid-level learning capability of CNNs.

430, **TITLE:** Grounding Referring Expressions in Images by Variational Context
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Grounding_Referring_Expressions_CVPR_2018_paper.html
AUTHORS: Hanwang Zhang, Yulei Niu, Shih-Fu Chang
HIGHLIGHT: In this paper, we propose a variational Bayesian method, called Variational Context, to solve the problem of complex context modeling in referring expression grounding.

431, **TITLE:** Dynamic Graph Generation Network: Generating Relational Knowledge From Diagrams
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_Dynamic_Graph_Generation_CVPR_2018_paper.html
AUTHORS: Daesik Kim, YoungJoon Yoo, Jee-Soo Kim, SangKuk Lee, Nojun Kwak
HIGHLIGHT: In this work, we introduce a new algorithm for analyzing a diagram, which contains visual and textual information in an abstract and integrated way.

432, **TITLE:** A Network Architecture for Point Cloud Classification via Automatic Depth Images Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Roveri_A_Network_Architecture_CVPR_2018_paper.html
AUTHORS: Riccardo Roveri, Lukas Rahmann, Cengiz Öztireli, Markus Gross
HIGHLIGHT: We propose a novel neural network architecture for point cloud classification.

433, **TITLE:** Towards Dense Object Tracking in a 2D Honeybee Hive
http://openaccess.thecvf.com/content_cvpr_2018/html/Bozek_Towards_Dense_Object_CVPR_2018_paper.html
AUTHORS: Katarzyna Bozek, Laetitia Hebert, Alexander S. Mikheyev, Greg J. Stephens
HIGHLIGHT: Given the novel application of CNNs in this study, we generate extensive problem-specific image data in which labeled examples are produced through a custom interface with Amazon Mechanical Turk.

434, **TITLE:** Long-Term On-Board Prediction of People in Traffic Scenes Under Uncertainty
http://openaccess.thecvf.com/content_cvpr_2018/html/Bhattacharyya_Long-Term_On-Board_Prediction_CVPR_2018_paper.html
AUTHORS: Apratim Bhattacharyya, Mario Fritz, Bernt Schiele
HIGHLIGHT: In this paper we argue that it is necessary to predict at least 1 second and we thus propose a new model that jointly predicts ego motion and people trajectories over such large time horizons.

435, **TITLE:** Single-Shot Refinement Neural Network for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Single-Shot_Refinement_Neural_CVPR_2018_paper.html
AUTHORS: Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, Stan Z. Li
HIGHLIGHT: To inherit the merits of both while overcoming their disadvantages, in this paper, we propose a novel single-shot based detector, called RefineDet, that achieves better accuracy than two-stage methods and maintains comparable efficiency of one-stage methods.

436, **TITLE:** Video Captioning via Hierarchical Reinforcement Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Video_Captioning_via_CVPR_2018_paper.html
AUTHORS: Xin Wang, Wenhui Chen, Jiawei Wu, Yuan-Fang Wang, William Yang Wang
HIGHLIGHT: This paper aims to address the challenge by proposing a novel hierarchical reinforcement learning framework for video captioning, where a high-level Manager module learns to design sub-goals and a low-level Worker module recognizes the primitive actions to fulfill the sub-goal.

437, **TITLE:** Tips and Tricks for Visual Question Answering: Learnings From the 2017 Challenge
http://openaccess.thecvf.com/content_cvpr_2018/html/Teney_Tips_and_Tricks_CVPR_2018_paper.html
AUTHORS: Damien Teney, Peter Anderson, Xiaodong He, Anton van den Hengel
HIGHLIGHT: This paper presents a state-of-the-art model for visual question answering (VQA), which won the first place in the 2017 VQA Challenge.

438, **TITLE:** Learning to Segment Every Thing
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Learning_to_Segment_CVPR_2018_paper.html
AUTHORS: Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, Ross Girshick
HIGHLIGHT: The goal of this paper is to propose a new partially supervised training paradigm, together with a novel weight transfer function, that enables training instance segmentation models on a large set of categories all of which have box annotations, but only a small fraction of which have mask annotations.

439, **TITLE:** Self-Supervised Adversarial Hashing Networks for Cross-Modal Retrieval

- http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Self-Supervised_Adversarial_Hashing_CVPR_2018_paper.html
AUTHORS: Chao Li, Cheng Deng, Ning Li, Wei Liu, Xinbo Gao, Dacheng Tao
HIGHLIGHT: In this paper, we propose a self-supervised adversarial hashing (SSAH) approach, which lies among the early attempts to incorporate adversarial learning into cross-modal hashing in a self-supervised fashion.
- 440, TITLE: Parallel Attention: A Unified Framework for Visual Object Discovery Through Dialogs and Queries
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhuang_Parallel_Attention_A_CVPR_2018_paper.html
AUTHORS: Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, Anton van den Hengel
HIGHLIGHT: In many of these cases natural language dialog is a natural way to specify the subject of interest, and the task achieving this capability (a.k.a, Referring Expression Comprehension) has recently attracted attention. To this end we propose a unified framework, the Parallel Attention (PLAN) network, to discover the object in an image that is being referred to in variable length natural expression descriptions, from short phrases query to long multi-round dialogs.
- 441, TITLE: Zigzag Learning for Weakly Supervised Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Zigzag_Learning_for_CVPR_2018_paper.html
AUTHORS: Xiaopeng Zhang, Jiashi Feng, Hongkai Xiong, Qi Tian
HIGHLIGHT: Unlike them, we propose a zigzag learning strategy to simultaneously discover reliable object instances and prevent the model from overfitting initial seeds.
- 442, TITLE: Attentive Fashion Grammar Network for Fashion Landmark Detection and Clothing Category Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Attentive_Fashion_Grammar_CVPR_2018_paper.html
AUTHORS: Wenguan Wang, Yuanlu Xu, Jianbing Shen, Song-Chun Zhu
HIGHLIGHT: This paper proposes a knowledge-guided fashion network to solve the problem of visual fashion analysis, e.g., fashion landmark localization and clothing category classification.
- 443, TITLE: Generalized Zero-Shot Learning via Synthesized Examples
http://openaccess.thecvf.com/content_cvpr_2018/html/Verma_Generalized_Zero-Shot_Learning_CVPR_2018_paper.html
AUTHORS: Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, Piyush Rai
HIGHLIGHT: We present a generative framework for generalized zero-shot learning where the training and test classes are not necessarily disjoint.
- 444, TITLE: Partially Shared Multi-Task Convolutional Neural Network With Local Constraint for Face Attribute Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Partially_Shared_Multi-Task_CVPR_2018_paper.html
AUTHORS: Jiajiong Cao, Yingming Li, Zhongfei Zhang
HIGHLIGHT: In this paper, we study the face attribute learning problem by considering the identity information and attribute relationships simultaneously.
- 445, TITLE: SYQ: Learning Symmetric Quantization for Efficient Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Faraone_SYQ_Learning_Symmetric_CVPR_2018_paper.html
AUTHORS: Julian Faraone, Nicholas Fraser, Michaela Blott, Philip H.W. Leong
HIGHLIGHT: In this paper, we introduce a quantization method to reduce this loss by learning a symmetric codebook for particular weight subgroups.
- 446, TITLE: DS*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems
http://openaccess.thecvf.com/content_cvpr_2018/html/Bernard_DS_Tighter_Lifting-Free_CVPR_2018_paper.html
AUTHORS: Florian Bernard, Christian Theobalt, Michael Moeller
HIGHLIGHT: In this work we study convex relaxations of quadratic optimisation problems over permutation matrices.
- 447, TITLE: Deep Mutual Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Deep_Mutual_Learning_CVPR_2018_paper.html
AUTHORS: Ying Zhang, Tao Xiang, Timothy M. Hospedales, Huchuan Lu
HIGHLIGHT: In this paper, we present a deep mutual learning (DML) strategy.
- 448, TITLE: Coupled End-to-End Transfer Learning With Generalized Fisher Information
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Coupled_End-to-End_Transfer_CVPR_2018_paper.html
AUTHORS: Shixing Chen, Caojin Zhang, Ming Dong
HIGHLIGHT: In this paper, we propose a novel Coupled End-to-end Transfer Learning (CETL) framework, which mainly consists of two convolutional neural networks (source and target) that connect to a shared decoder.
- 449, TITLE: Residual Parameter Transfer for Deep Domain Adaptation

- http://openaccess.thecvf.com/content_cvpr_2018/html/Rozantsev_Residual_Parameter_Transfer_CVPR_2018_paper.html
AUTHORS: Artem Rozantsev, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: By contrast, we introduce a network architecture that includes auxiliary residual networks, which we train to predict the parameters in the domain with little annotated data from those in the other one.
- 450, TITLE: High-Order Tensor Regularization With Application to Attribute Ranking
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_High-Order_Tensor_Regularization_CVPR_2018_paper.html
AUTHORS: Kwang In Kim, Juhyun Park, James Tompkin
HIGHLIGHT: Our new method for intrinsically regularizing and learning tensors on Riemannian manifolds introduces a surrogate object to encapsulate the geometric characteristic of the tensor.
- 451, TITLE: Learning to Localize Sound Source in Visual Scenes
http://openaccess.thecvf.com/content_cvpr_2018/html/Senocak_Learning_to_Localize_CVPR_2018_paper.html
AUTHORS: Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, In So Kweon
HIGHLIGHT: In this paper, we propose a novel unsupervised algorithm to address the problem of localizing the sound source in visual scenes.
- 452, TITLE: Dynamic Few-Shot Visual Learning Without Forgetting
http://openaccess.thecvf.com/content_cvpr_2018/html/Gidaris_Dynamic_Few-Shot_Visual_CVPR_2018_paper.html
AUTHORS: Spyros Gidaris, Nikos Komodakis
HIGHLIGHT: In this context, the goal of our work is to devise a few-shot visual learning system that during test time it will be able to efficiently learn novel categories from only a few training data while at the same time it will not forget the initial categories on which it was trained (here called base categories).
- 453, TITLE: Two-Step Quantization for Low-Bit Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Two-Step_Quantization_for_CVPR_2018_paper.html
AUTHORS: Peisong Wang, Qinghao Hu, Yifan Zhang, Chunjie Zhang, Yang Liu, Jian Cheng
HIGHLIGHT: In this paper, we propose a simple yet effective Two-Step Quantization (TSQ) framework, by decomposing the network quantization problem into two steps: code learning and transformation function learning based on the learned codes.
- 454, TITLE: Improved Lossy Image Compression With Priming and Spatially Adaptive Bit Rates for Recurrent Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Johnston_Improved_Lossy_Image_CVPR_2018_paper.html
AUTHORS: Nick Johnston, Damien Vincent, David Minnen, Michele Covell, Saurabh Singh, Troy Chinen, Sung Jin Hwang, Joel Shor, George Toderici
HIGHLIGHT: We propose a method for lossy image compression based on recurrent, convolutional neural networks that outperforms BPG (4:2:0), WebP, JPEG2000, and JPEG as measured by MS-SSIM.
- 455, TITLE: Conditional Probability Models for Deep Image Compression
http://openaccess.thecvf.com/content_cvpr_2018/html/Mentzer_Conditional_Probability_Models_CVPR_2018_paper.html
AUTHORS: Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, Luc Van Gool
HIGHLIGHT: In this paper, we focus on the latter challenge and propose a new technique to navigate the rate-distortion trade-off for an image compression auto-encoder.
- 456, TITLE: Deep Diffeomorphic Transformer Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Detlefsen_Deep_Diffeomorphic_Transformer_CVPR_2018_paper.html
AUTHORS: Nicki Skafte Detlefsen, Oren Freifeld, Søren Hauberg
HIGHLIGHT: We investigate the use of flexible diffeomorphic image transformations within such networks and demonstrate that significant performance gains can be attained over currently-used models.
- 457, TITLE: The Lovász-Softmax Loss: A Tractable Surrogate for the Optimization of the Intersection-Over-Union Measure in Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Berman_The_Lovasz-Softmax_Loss_CVPR_2018_paper.html
AUTHORS: Maxim Berman, Amal Rannen Triki, Matthew B. Blaschko
HIGHLIGHT: We present a method for direct optimization of the mean intersection-over-union loss in neural networks, in the context of semantic image segmentation, based on the convex Lovász extension of submodular losses.
- 458, TITLE: Generative Adversarial Perturbations
http://openaccess.thecvf.com/content_cvpr_2018/html/Poursaeed_Generative_Adversarial_Perturbations_CVPR_2018_paper.html
AUTHORS: Omid Poursaeed, Isay Katsman, Bicheng Gao, Serge Belongie
HIGHLIGHT: In this paper, we propose novel generative models for creating adversarial examples, slightly perturbed images resembling natural images but maliciously crafted to fool pre-trained models.

- 459, TITLE: Learning Strict Identity Mappings in Deep Residual Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Learning_Strict_Identity_CVPR_2018_paper.html
AUTHORS: Xin Yu, Zhiding Yu, Srikumar Ramalingam
HIGHLIGHT: In this paper, we propose ϵ -ResNet that allows us to automatically discard redundant layers, which produces responses that are smaller than a threshold ϵ , without any loss in performance.
- 460, TITLE: Geometric Robustness of Deep Networks: Analysis and Improvement
http://openaccess.thecvf.com/content_cvpr_2018/html/Kanbak_Geometric_Robustness_of_CVPR_2018_paper.html
AUTHORS: Can Kanbak, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard
HIGHLIGHT: We propose ManiFool as a simple yet scalable algorithm to measure the invariance of deep networks.
- 461, TITLE: View Extrapolation of Human Body From a Single Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_View_Extrapolation_of_CVPR_2018_paper.html
AUTHORS: Hao Zhu, Hao Su, Peng Wang, Xun Cao, Ruigang Yang
HIGHLIGHT: To address the problem, we propose a novel deep learning based pipeline that explicitly estimates and leverages the geometry of the underlying human body.
- 462, TITLE: Geometry Aware Constrained Optimization Techniques for Deep Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Roy_Geometry_Aware_Constrained_CVPR_2018_paper.html
AUTHORS: Soumava Kumar Roy, Zakaria Mhammedi, Mehrtash Harandi
HIGHLIGHT: In this paper, we generalize the Stochastic Gradient Descent (SGD) and RMSProp algorithms to the setting of Riemannian optimization.
- 463, TITLE: PointNetVLAD: Deep Point Cloud Based Retrieval for Large-Scale Place Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Uy_PointNetVLAD_Deep_Point_CVPR_2018_paper.html
AUTHORS: Mikaela Angelina Uy, Gim Hee Lee
HIGHLIGHT: In this paper, we propose the PointNetVLAD where we leverage on the recent success of deep networks to solve point cloud based retrieval for place recognition.
We create benchmark datasets for point cloud based retrieval for place recognition, and the experimental results on these datasets show the feasibility of our PointNetVLAD.
- 464, TITLE: An Efficient and Provable Approach for Mixture Proportion Estimation Using Linear Independence Assumption
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_An_Efficient_and_CVPR_2018_paper.html
AUTHORS: Xiyu Yu, Tongliang Liu, Mingming Gong, Kayhan Batmanghelich, Dacheng Tao
HIGHLIGHT: In this paper, we study the mixture proportion estimation (MPE) problem in a new setting: given samples from the mixture and the component distributions, we identify the proportions of the components in the mixture distribution.
- 465, TITLE: VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_VoxelNet_End-to-End_Learning_CVPR_2018_paper.html
AUTHORS: Yin Zhou, Oncel Tuzel
HIGHLIGHT: In this work, we remove the need of manual feature engineering for 3D point clouds and propose VoxelNet, a generic 3D detection network that unifies feature extraction and bounding box prediction into a single stage, end-to-end trainable deep network.
- 466, TITLE: Image to Image Translation for Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Murez_Image_to_Image_CVPR_2018_paper.html
AUTHORS: Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, Kyungnam Kim
HIGHLIGHT: We propose a general framework for unsupervised domain adaptation, which allows deep neural networks trained on a source domain to be tested on a different target domain without requiring any training annotations in the target domain.
- 467, TITLE: MobileNetV2: Inverted Residuals and Linear Bottlenecks
http://openaccess.thecvf.com/content_cvpr_2018/html/Sandler_MobileNetV2_Inverted_Residuals_CVPR_2018_paper.html
AUTHORS: Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, Liang-Chieh Chen
HIGHLIGHT: In this paper we describe a new mobile architecture, $mbox$ {MobileNetV2}, that improves the state of the art performance of mobile models on multiple tasks and benchmarks as well as across a spectrum of different model sizes.
- 468, TITLE: Im2Struct: Recovering 3D Shape Structure From a Single RGB Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Niu_Im2Struct_Recovering_3D_CVPR_2018_paper.html

- AUTHORS: Chengjie Niu, Jun Li, Kai Xu
HIGHLIGHT: We propose to recover 3D shape structures from single RGB images, where structure refers to shape parts represented by cuboids and part relations encompassing connectivity and symmetry.
- 469, TITLE: Trust Your Model: Light Field Depth Estimation With Inline Occlusion Handling
http://openaccess.thecvf.com/content_cvpr_2018/html/Schilling_Trust_Your_Model_CVPR_2018_paper.html
AUTHORS: Hendrik Schilling, Maximilian Diebold, Carsten Rother, Bernd Jähne
HIGHLIGHT: Our main contribution is a new way to handle occlusions which improves general accuracy and quality of object borders.
- 470, TITLE: Baseline Desensitizing in Translation Averaging
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhuang_Baseline_Desensitizing_in_CVPR_2018_paper.html
AUTHORS: Bingbing Zhuang, Loong-Fah Cheong, Gim Hee Lee
HIGHLIGHT: In this paper, we carefully design a simple yet effective bilinear objective function, introducing a variable to perform the requisite normalization.
- 471, TITLE: Mining Point Cloud Local Structures by Kernel Correlation and Graph Pooling
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Mining_Point_Cloud_CVPR_2018_paper.html
AUTHORS: Yiru Shen, Chen Feng, Yaoqing Yang, Dong Tian
HIGHLIGHT: In this regard, we present two new operations to improve PointNet with a more efficient exploitation of local structures.
- 472, TITLE: Large-Scale Point Cloud Semantic Segmentation With Superpoint Graphs
http://openaccess.thecvf.com/content_cvpr_2018/html/Landrieu_Large-Scale_Point_Cloud_CVPR_2018_paper.html
AUTHORS: Loic Landrieu, Martin Simonovsky
HIGHLIGHT: We propose a novel deep learning-based framework to tackle the challenge of semantic segmentation of large-scale point clouds of millions of points.
- 473, TITLE: Very Large-Scale Global SfM by Distributed Motion Averaging
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_Very_Large-Scale_Global_CVPR_2018_paper.html
AUTHORS: Siyu Zhu, Runze Zhang, Lei Zhou, Tianwei Shen, Tian Fang, Ping Tan, Long Quan
HIGHLIGHT: This work proposes a divide-and-conquer framework to solve very large global SfM at the scale of millions of images.
- 474, TITLE: ScanComplete: Large-Scale Scene Completion and Semantic Segmentation for 3D Scans
http://openaccess.thecvf.com/content_cvpr_2018/html/Dai_ScanComplete_Large-Scale_Scene_CVPR_2018_paper.html
AUTHORS: Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, Matthias Nießner
HIGHLIGHT: We introduce ScanComplete, a novel data-driven approach for taking an incomplete 3D scan of a scene as input and predicting a complete 3D model along with per-voxel semantic labels.
- 475, TITLE: Solving the Perspective-2-Point Problem for Flying-Camera Photo Composition
http://openaccess.thecvf.com/content_cvpr_2018/html/Lan_Solving_the_Perspective-2-Point_CVPR_2018_paper.html
AUTHORS: Ziquan Lan, David Hsu, Gim Hee Lee
HIGHLIGHT: We model it as a Perspective-2-Point (P2P) problem, which is under-constrained to determine the six degrees-of-freedom camera pose uniquely.
- 476, TITLE: Reflection Removal for Large-Scale 3D Point Clouds
http://openaccess.thecvf.com/content_cvpr_2018/html/Yun_Reflection_Removal_for_CVPR_2018_paper.html
AUTHORS: Jae-Seong Yun, Jae-Young Sim
HIGHLIGHT: In this paper, we propose an efficient reflection removal algorithm for LS3DPCs.
- 477, TITLE: Attentional ShapeContextNet for Point Cloud Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Xie_Attentional_ShapeContextNet_for_CVPR_2018_paper.html
AUTHORS: Saining Xie, Sainan Liu, Zeyu Chen, Zhuowen Tu
HIGHLIGHT: We tackle the problem of point cloud recognition.
- 478, TITLE: Geometry-Aware Deep Network for Single-Image Novel View Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Geometry-Aware_Deep_Network_CVPR_2018_paper.html
AUTHORS: Miaomiao Liu, Xuming He, Mathieu Salzmann
HIGHLIGHT: By contrast, in this paper, we propose to exploit the 3D geometry of the scene to synthesize a novel view.

479, TITLE: InverseFaceNet: Deep Monocular Inverse Face Rendering
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_InverseFaceNet_Deep_Monocular_CVPR_2018_paper.html
AUTHORS: Hyeonwoo Kim, Michael Zollhöfer, Ayush Tewari, Justus Thies, Christian Richardt, Christian Theobalt
HIGHLIGHT: We introduce InverseFaceNet, a deep convolutional inverse rendering framework for faces that jointly estimates facial pose, shape, expression, reflectance and illumination from a single input image.

480, TITLE: Sparse Photometric 3D Face Reconstruction Guided by Morphable Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Sparse_Photometric_3D_CVPR_2018_paper.html
AUTHORS: Xuan Cao, Zhang Chen, Anpei Chen, Xin Chen, Shiyong Li, Jingyi Yu
HIGHLIGHT: We present a novel 3D face reconstruction technique that leverages sparse photometric stereo (PS) and latest advances on face registration / modeling from a single image.

481, TITLE: Texture Mapping for 3D Reconstruction With RGB-D Sensor
http://openaccess.thecvf.com/content_cvpr_2018/html/Fu_Texture_Mapping_for_CVPR_2018_paper.html
AUTHORS: Yanping Fu, Qingan Yan, Long Yang, Jie Liao, Chunxia Xiao
HIGHLIGHT: In this paper, we propose a global-to-local correction strategy to obtain more desired texture mapping results.

482, TITLE: Learning Less Is More - 6D Camera Localization via 3D Surface Regression
http://openaccess.thecvf.com/content_cvpr_2018/html/Brachmann_Learning_Less_Is_CVPR_2018_paper.html
AUTHORS: Eric Brachmann, Carsten Rother
HIGHLIGHT: In this work, we address the task of predicting the 6D camera pose from a single RGB image in a given 3D environment.

483, TITLE: Feature Mapping for Learning Fast and Accurate 3D Pose Inference From Synthetic Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Rad_Feature_Mapping_for_CVPR_2018_paper.html
AUTHORS: Mahdi Rad, Markus Oberweger, Vincent Lepetit
HIGHLIGHT: We propose a simple and efficient method for exploiting synthetic images when training a Deep Network to predict a 3D pose from an image.

484, TITLE: Indoor RGB-D Compass From a Single Line and Plane
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_Indoor_RGB-D_Compass_CVPR_2018_paper.html
AUTHORS: Pyojin Kim, Brian Coltin, H. Jin Kim
HIGHLIGHT: We propose a novel approach to estimate the three degrees of freedom (DoF) drift-free rotational motion of an RGB-D camera from only a single line and plane in the Manhattan world (MW).

485, TITLE: Geometry-Aware Network for Non-Rigid Shape Prediction From a Single View
http://openaccess.thecvf.com/content_cvpr_2018/html/Pumarola_Geometry-Aware_Network_for_CVPR_2018_paper.html
AUTHORS: Albert Pumarola, Antonio Agudo, Lorenzo Porzi, Alberto Sanfeliu, Vincent Lepetit, Francesc Moreno-Noguer
HIGHLIGHT: We propose a method for predicting the 3D shape of a deformable surface from a single view.

486, TITLE: Sim2Real Viewpoint Invariant Visual Servoing by Recurrent Control
http://openaccess.thecvf.com/content_cvpr_2018/html/Sadeghi_Sim2Real_Viewpoint_Invariant_CVPR_2018_paper.html
AUTHORS: Fereshteh Sadeghi, Alexander Toshev, Eric Jang, Sergey Levine
HIGHLIGHT: In this paper, we propose learning viewpoint invariant visual servoing skills in a robot manipulation task.

487, TITLE: DocUNet: Document Image Unwarping via a Stacked U-Net
http://openaccess.thecvf.com/content_cvpr_2018/html/Ma_DocUNet_Document_Image_CVPR_2018_paper.html
AUTHORS: Ke Ma, Zhixin Shu, Xue Bai, Jue Wang, Dimitris Samaras
HIGHLIGHT: In this paper, we develop the first learning-based method to achieve this goal. Because large-scale real-world data with ground truth deformation is difficult to obtain, we create a synthetic dataset with approximately 100 thousand images by warping non-distorted document images. We further create a comprehensive benchmark that covers various real-world conditions.

488, TITLE: Analysis of Hand Segmentation in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Urooj_Analysis_of_Hand_CVPR_2018_paper.html
AUTHORS: Aisha Urooj, Ali Borji
HIGHLIGHT: We fine-tune RefineNet, a leading semantic segmentation method, for hand segmentation and find that it does much better than the best contenders.

Finally, we annotate a subset of the EgoHands dataset for fine-grained action recognition and show that an accuracy of 58.6% can be achieved by just looking at a single hand pose which is much better than the chance level (12.5%).

- 489, TITLE: RoadTracer: Automatic Extraction of Road Networks From Aerial Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Bastani_RoadTracer_Automatic_Extraction_CVPR_2018_paper.html
AUTHORS: Favyen Bastani, Songtao He, Sofiane Abbar, Mohammad Alizadeh, Hari Balakrishnan, Sanjay Chawla, Sam Madden, David DeWitt
HIGHLIGHT: We propose RoadTracer, a new method to automatically construct accurate road network maps from aerial images.
- 490, TITLE: Alternating-Stereo VINS: Observability Analysis and Performance Evaluation
http://openaccess.thecvf.com/content_cvpr_2018/html/Paul_Alternating-Stereo_VINS_Observability_CVPR_2018_paper.html
AUTHORS: Mrinal K. Paul, Stergios I. Roumeliotis
HIGHLIGHT: To address this limitation, in this work, a novel two-camera alternating-stereo VINS is presented.
- 491, TITLE: Soccer on Your Tabletop
http://openaccess.thecvf.com/content_cvpr_2018/html/Rematas_Soccer_on_Your_CVPR_2018_paper.html
AUTHORS: Konstantinos Rematas, Ira Kemelmacher-Shlizerman, Brian Curless, Steve Seitz
HIGHLIGHT: We present a system that transforms a monocular video of a soccer game into a moving 3D reconstruction, in which the players and field can be rendered interactively with a 3D viewer or through an Augmented Reality device.
- 492, TITLE: EPINET: A Fully-Convolutional Neural Network Using Epipolar Geometry for Depth From Light Field Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Shin_EPINET_A_Fully-Convolutional_CVPR_2018_paper.html
AUTHORS: Changha Shin, Hae-Gon Jeon, Youngjin Yoon, In So Kweon, Seon Joo Kim
HIGHLIGHT: In this paper, we introduce a fast and accurate light field depth estimation method based on a fully-convolutional neural network.
- 493, TITLE: A Hybrid l1-l0 Layer Decomposition Model for Tone Mapping
http://openaccess.thecvf.com/content_cvpr_2018/html/Liang_A_Hybrid_l1-l0_CVPR_2018_paper.html
AUTHORS: Zhetong Liang, Jun Xu, David Zhang, Zisheng Cao, Lei Zhang
HIGHLIGHT: In this paper, we propose a hybrid L1-L0 decomposition model to address these problems.
- 494, TITLE: Deeply Learned Filter Response Functions for Hyperspectral Reconstruction
http://openaccess.thecvf.com/content_cvpr_2018/html/Nie_Deeply_Learned_Filter_CVPR_2018_paper.html
AUTHORS: Shijie Nie, Lin Gu, Yinqiang Zheng, Antony Lam, Nobutaka Ono, Imari Sato
HIGHLIGHT: In this paper, rather than use RGB spectral responses, we simultaneously learn optimized camera spectral response functions (to be implemented in hardware) and a mapping for spectral reconstruction by using an end-to-end network.
- 495, TITLE: CRRN: Multi-Scale Guided Concurrent Reflection Removal Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Wan_CRRN_Multi-Scale_Guided_CVPR_2018_paper.html
AUTHORS: Renjie Wan, Boxin Shi, Ling-Yu Duan, Ah-Hwee Tan, Alex C. Kot
HIGHLIGHT: In this paper, we propose the Concurrent Reflection Removal Network (CRRN) to tackle this problem in a unified framework.
- 496, TITLE: Single Image Reflection Separation With Perceptual Losses
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Single_Image_Reflection_CVPR_2018_paper.html
AUTHORS: Xuaner Zhang, Ren Ng, Qifeng Chen
HIGHLIGHT: We present an approach to separating reflection from a single image.
We create a dataset of real-world images with reflection and corresponding ground-truth transmission layers for quantitative evaluation and model training.
- 497, TITLE: A Robust Method for Strong Rolling Shutter Effects Correction Using Lines With Automatic Feature Selection
http://openaccess.thecvf.com/content_cvpr_2018/html/Lao_A_Robust_Method_CVPR_2018_paper.html
AUTHORS: Yizhen Lao, Omar Ait-Aider
HIGHLIGHT: We present a robust method which compensates RS distortions in a single image using a set of image curves, basing on the knowledge that they correspond to 3D straight lines.
- 498, TITLE: Time-Resolved Light Transport Decomposition for Thermal Photometric Stereo
http://openaccess.thecvf.com/content_cvpr_2018/html/Tanaka_Time-Resolved_Light_Transport_CVPR_2018_paper.html

AUTHORS: Kenichiro Tanaka, Nobuhiro Ikeya, Tsuyoshi Takatani, Hiroyuki Kubo, Takuya Funatomi, Yasuhiro Mukaigawa
HIGHLIGHT: We present a novel time-resolved light transport decomposition method using thermal imaging.

499, TITLE: Efficient Diverse Ensemble for Discriminative Co-Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Meshgi_Efficient_Diverse_Ensemble_CVPR_2018_paper.html
AUTHORS: Kouros Meshgi, Shigeyuki Oba, Shin Ishii
HIGHLIGHT: In this study, we propose an online ensemble tracker that directly generates a diverse committee by generating an efficient set of artificial training.

500, TITLE: Rolling Shutter and Radial Distortion Are Features for High Frame Rate Multi-Camera Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Bapat_Rolling_Shutter_and_CVPR_2018_paper.html
AUTHORS: Akash Bapat, True Price, Jan-Michael Frahm
HIGHLIGHT: In this paper, we introduce a novel multi-camera tracking approach that for the first time jointly leverages the information introduced by rolling shutter and radial distortion as a feature to achieve superior performance with respect to high-frequency camera pose estimation.

501, TITLE: A Twofold Siamese Network for Real-Time Object Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/He_A_Twofold_Siamese_CVPR_2018_paper.html
AUTHORS: Anfeng He, Chong Luo, Xinmei Tian, Wenjun Zeng
HIGHLIGHT: Observing that Semantic features learned in an image classification task and Appearance features learned in a similarity matching task complement each other, we build a twofold Siamese network, named SA-Siam, for real-time object tracking.

502, TITLE: Multi-Cue Correlation Filters for Robust Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Multi-Cue_Correlation_Filters_CVPR_2018_paper.html
AUTHORS: Ning Wang, Wengang Zhou, Qi Tian, Richang Hong, Meng Wang, Houqiang Li
HIGHLIGHT: In this paper, we propose an efficient multi-cue analysis framework for robust visual tracking.

503, TITLE: Learning Attentions: Residual Attentional Siamese Network for High Performance Online Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Learning_Attentions_Residual_CVPR_2018_paper.html
AUTHORS: Qiang Wang, Zhu Teng, Junliang Xing, Jin Gao, Weiming Hu, Stephen Maybank
HIGHLIGHT: This work presents a Residual Attentional Siamese Network (RASNet) for high performance object tracking.

504, TITLE: SINT++: Robust Visual Tracking via Adversarial Positive Instance Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_SINT_Robust_Visual_CVPR_2018_paper.html
AUTHORS: Xiao Wang, Chenglong Li, Bin Luo, Jin Tang
HIGHLIGHT: In this paper, we propose to generate hard positive samples via adversarial learning for visual tracking.

505, TITLE: High-Speed Tracking With Multi-Kernel Correlation Filters
http://openaccess.thecvf.com/content_cvpr_2018/html/Tang_High-Speed_Tracking_With_CVPR_2018_paper.html
AUTHORS: Ming Tang, Bin Yu, Fan Zhang, Jinqiao Wang
HIGHLIGHT: In this paper, we will introduce the MKL into KCF in a different way than MKCF.

506, TITLE: Occlusion Aware Unsupervised Learning of Optical Flow
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Occlusion_Aware_Unsupervised_CVPR_2018_paper.html
AUTHORS: Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, Peng Wang, Wei Xu
HIGHLIGHT: In this work we introduce a new method which models occlusion explicitly and a new warping way that facilitates the learning of large motion.

507, TITLE: Revisiting Video Saliency: A Large-Scale Benchmark and a New Model
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Revisiting_Video_Saliency_CVPR_2018_paper.html
AUTHORS: Wenguan Wang, Jianbing Shen, Fang Guo, Ming-Ming Cheng, Ali Borji
HIGHLIGHT: In this work, we contribute to video saliency research in two ways. First, we introduce a new benchmark for predicting human eye movements during dynamic scene free-viewing, which is long-time urged in this field.

508, TITLE: Learning Spatial-Temporal Regularized Correlation Filters for Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Learning_Spatial-Temporal-Regularized_CVPR_2018_paper.html
AUTHORS: Feng Li, Cheng Tian, Wangmeng Zuo, Lei Zhang, Ming-Hsuan Yang

HIGHLIGHT: In this work, by introducing temporal regularization to SRDCF with single sample, we present our spatial-temporal regularized correlation filters (STRCF).

509, **TITLE:** Multimodal Visual Concept Learning With Weakly Supervised Techniques
http://openaccess.thecvf.com/content_cvpr_2018/html/Bouritsas_Multimodal_Visual_Concept_CVPR_2018_paper.html
AUTHORS: Giorgos Bouritsas, Petros Koutras, Athanasia Zlatintsi, Petros Maragos
HIGHLIGHT: Towards this goal, in this paper we use textual cues as means of supervision, introducing two weakly supervised techniques that extend the Multiple Instance Learning (MIL) framework: the Fuzzy Sets Multiple Instance Learning (FSMIL) and the Probabilistic Labels Multiple Instance Learning (PLMIL).

510, **TITLE:** Efficient Large-Scale Approximate Nearest Neighbor Search on OpenCL FPGA
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Efficient_Large-Scale_Approximate_CVPR_2018_paper.html
AUTHORS: Jiali Zhang, Soroosh Khoram, Jing Li
HIGHLIGHT: We present a new method for Product Quantization (PQ) based approximated nearest neighbor search (ANN) in high dimensional spaces.

511, **TITLE:** Learning a Complete Image Indexing Pipeline
http://openaccess.thecvf.com/content_cvpr_2018/html/Jain_Learning_a_Complete_CVPR_2018_paper.html
AUTHORS: Himalaya Jain, Joaquin Zepeda, Patrick Pérez, Rémi Gribonval
HIGHLIGHT: In this work, we propose a first system that learns both components within a unifying neural framework of structured binary encoding.

512, **TITLE:** Transparency by Design: Closing the Gap Between Performance and Interpretability in Visual Reasoning
http://openaccess.thecvf.com/content_cvpr_2018/html/Mascharka_Transparency_by_Design_CVPR_2018_paper.html
AUTHORS: David Mascharka, Philip Tran, Ryan Soklaski, Arjun Majumdar
HIGHLIGHT: In this paper, we close the performance gap between interpretable models and state-of-the-art visual reasoning methods.
We propose a set of visual-reasoning primitives which, when composed, manifest as a model capable of performing complex reasoning tasks in an explicitly-interpretable manner.

513, **TITLE:** Fooling Vision and Language Models Despite Localization and Attention Mechanism
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Fooling_Vision_and_CVPR_2018_paper.html
AUTHORS: Xiaojun Xu, Xinyun Chen, Chang Liu, Anna Rohrbach, Trevor Darrell, Dawn Song
HIGHLIGHT: In this paper, we study adversarial examples for vision and language models, which incorporate natural language understanding and complex structures such as attention, localization, and modular architectures.

514, **TITLE:** Categorizing Concepts With Basic Level for Vision-to-Language
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Categorizing_Concepts_With_CVPR_2018_paper.html
AUTHORS: Hanzhang Wang, Hanli Wang, Kaisheng Xu
HIGHLIGHT: Inspired by the basic level in early cognition, a Basic Concept (BaC) category is proposed in this work that contains both consensus and proper level of visual content to help neural network tackle the above problems.

515, **TITLE:** Don't Just Assume; Look and Answer: Overcoming Priors for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Agrawal_Dont_Just_Assume_CVPR_2018_paper.html
AUTHORS: Aishwarya Agrawal, Dhruv Batra, Devi Parikh, Aniruddha Kembhavi
HIGHLIGHT: To encourage development of models geared towards the latter, we propose a new setting for VQA where for every question type, train and test sets have different prior distributions of answers.

516, **TITLE:** Learning Pixel-Level Semantic Affinity With Image-Level Supervision for Weakly Supervised Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Ahn_Learning_Pixel-Level_Semantic_CVPR_2018_paper.html
AUTHORS: Jiwoon Ahn, Suha Kwak
HIGHLIGHT: To alleviate this issue, we present a novel framework that generates segmentation labels of images given their image-level class labels.

517, **TITLE:** From Lifestyle Vlogs to Everyday Interactions
http://openaccess.thecvf.com/content_cvpr_2018/html/Fouhey_From_Lifestyle_Vlogs_CVPR_2018_paper.html
AUTHORS: David F. Fouhey, Wei-cheng Kuo, Alexei A. Efros, Jitendra Malik
HIGHLIGHT: In this work, we do the reverse and search implicitly: we start with a large collection of interaction-rich video data and then annotate and analyze it.

518, TITLE: Cross-Domain Weakly-Supervised Object Detection Through Progressive Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Inoue_Cross-Domain_Weakly-Supervised_Object_CVPR_2018_paper.html
AUTHORS: Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, Kiyoharu Aizawa
HIGHLIGHT: In this paper, we present a framework for a novel task, cross-domain weakly supervised object detection, which addresses this question.

519, TITLE: RotationNet: Joint Object Categorization and Pose Estimation Using Multiviews From Unsupervised Viewpoints
http://openaccess.thecvf.com/content_cvpr_2018/html/Kanezaki_RotationNet_Joint_Object_CVPR_2018_paper.html
AUTHORS: Asako Kanezaki, Yasuyuki Matsushita, Yoshifumi Nishida
HIGHLIGHT: We propose a Convolutional Neural Network (CNN)-based model "RotationNet," which takes multi-view images of an object as input and jointly estimates its pose and object category.

520, TITLE: An End-to-End TextSpotter With Explicit Alignment and Attention
http://openaccess.thecvf.com/content_cvpr_2018/html/He_An_End-to-End_TextSpotter_CVPR_2018_paper.html
AUTHORS: Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, Changming Sun
HIGHLIGHT: In this work, we present a conceptually simple yet efficient framework that simultaneously processes the two tasks in a united framework.

521, TITLE: WILDTRACK: A Multi-Camera HD Dataset for Dense Unscripted Pedestrian Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Chavdarova_WILDTRACK_A_Multi-Camera_CVPR_2018_paper.html
AUTHORS: Tatjana Chavdarova, Pierre Baqué, Stéphane Bouquet, Andrii Maksai, Cijo Jose, Timur Bagautdinov, Louis Letry, Pascal Fua, Luc Van Gool, François Fleuret
HIGHLIGHT: In this paper, we present a new large-scale and high-resolution dataset.

522, TITLE: Direct Shape Regression Networks for End-to-End Face Alignment
http://openaccess.thecvf.com/content_cvpr_2018/html/Miao_Direct_Shape_Regression_CVPR_2018_paper.html
AUTHORS: Xin Miao, Xiantong Zhen, Xianglong Liu, Cheng Deng, Vassilis Athitsos, Heng Huang
HIGHLIGHT: In this paper, we propose the direct shape regression network (DSRN) for end-to-end face alignment by jointly handling the aforementioned challenges in a unified framework.

523, TITLE: Natural and Effective Obfuscation by Head Inpainting
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Natural_and_Effective_CVPR_2018_paper.html
AUTHORS: Qianru Sun, Liqian Ma, Seong Joon Oh, Luc Van Gool, Bernt Schiele, Mario Fritz
HIGHLIGHT: In this work, we propose a novel head inpainting obfuscation technique.

524, TITLE: 3D Semantic Trajectory Reconstruction From 3D Pixel Continuum
http://openaccess.thecvf.com/content_cvpr_2018/html/Yoon_3D_Semantic_Trajectory_CVPR_2018_paper.html
AUTHORS: Jae Shin Yoon, Ziwei Li, Hyun Soo Park
HIGHLIGHT: This paper presents a method to reconstruct dense semantic trajectory stream of human interactions in 3D from synchronized multiple videos.

525, TITLE: Optimizing Filter Size in Convolutional Neural Networks for Facial Action Unit Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Han_Optimizing_Filter_Size_CVPR_2018_paper.html
AUTHORS: Shizhong Han, Zibo Meng, Zhiyuan Li, James O'Reilly, Jie Cai, Xiaofeng Wang, Yan Tong
HIGHLIGHT: This paper proposes a novel Optimized Filter Size CNN (OFS-CNN), where the filter sizes and weights of all convolutional layers are learned simultaneously from the training data along with learning convolution filters.

526, TITLE: V2V-PoseNet: Voxel-to-Voxel Prediction Network for Accurate 3D Hand and Human Pose Estimation From a Single Depth Map
http://openaccess.thecvf.com/content_cvpr_2018/html/Moon_V2V-PoseNet_Voxel-to-Voxel_Prediction_CVPR_2018_paper.html
AUTHORS: Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee
HIGHLIGHT: We design our model as a 3D CNN that provides accurate estimates while running in real-time.

527, TITLE: Ring Loss: Convex Feature Normalization for Face Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zheng_Ring_Loss_Convex_CVPR_2018_paper.html
AUTHORS: Yutong Zheng, Dipan K. Pal, Marios Savvides
HIGHLIGHT: We apply Ring loss to large-scale face recognition problems and present results on LFW, the challenging protocols of IJB-A Janus, Janus CS3 (a superset of IJB-A Janus), Celebrity Frontal-Profile (CFP) and MegaFace with 1 million distractors.

- 528, TITLE: Adversarially Occluded Samples for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Adversarially_Occluded_Samples_CVPR_2018_paper.html
AUTHORS: Houjing Huang, Dangwei Li, Zhang Zhang, Xiaotang Chen, Kaiqi Huang
HIGHLIGHT: Considering this fact, we propose to augment the variation of training data by introducing Adversarially Occluded Samples.
- 529, TITLE: Classifier Learning With Prior Probabilities for Facial Action Unit Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Classifier_Learning_With_CVPR_2018_paper.html
AUTHORS: Yong Zhang, Weiming Dong, Bao-Gang Hu, Qiang Ji
HIGHLIGHT: To alleviate this issue, we propose a knowledge-driven method for jointly learning multiple AU classifiers without any AU annotation by leveraging prior probabilities on AUs, including expression-independent and expression-dependent AU probabilities.
- 530, TITLE: 4DFAB: A Large Scale 4D Database for Facial Expression Analysis and Biometric Applications
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_4DFAB_A_Large_CVPR_2018_paper.html
AUTHORS: Shiyang Cheng, Irene Kotsia, Maja Pantic, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we conduct several experiments and demonstrate the usefulness of the database in various applications.
- 531, TITLE: Seeing Small Faces From Robust Anchor's Perspective
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_Seeing_Small_Faces_CVPR_2018_paper.html
AUTHORS: Chenchen Zhu, Ran Tao, Khoa Luu, Marios Savvides
HIGHLIGHT: In this paper, we investigate why this is the case.
- 532, TITLE: 2D/3D Pose Estimation and Action Recognition Using Multitask Deep Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Luvizon_2D3D_Pose_Estimation_CVPR_2018_paper.html
AUTHORS: Diogo C. Luvizon, David Picard, Hedi Tabia
HIGHLIGHT: In this work, we propose a multitask framework for jointly 2D and 3D pose estimation from still images and human action recognition from video sequences.
- 533, TITLE: Dense 3D Regression for Hand Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wan_Dense_3D_Regression_CVPR_2018_paper.html
AUTHORS: Chengde Wan, Thomas Probst, Luc Van Gool, Angela Yao
HIGHLIGHT: We present a simple and effective method for 3D hand pose estimation from a single depth frame.
- 534, TITLE: Camera Style Adaptation for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhong_Camera_Style_Adaptation_CVPR_2018_paper.html
AUTHORS: Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, Yi Yang
HIGHLIGHT: In this paper, we explicitly consider this challenge by introducing camera style (CamStyle) adaptation.
- 535, TITLE: PoseTrack: A Benchmark for Human Pose Estimation and Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Andriluka_PoseTrack_A_Benchmark_CVPR_2018_paper.html
AUTHORS: Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, Bernt Schiele
HIGHLIGHT: To address this shortcoming this paper introduces PoseTrack which is a new large-scale benchmark for video-based human pose estimation and articulated tracking.
- 536, TITLE: Exploit the Unknown Gradually: One-Shot Video-Based Person Re-Identification by Stepwise Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Exploit_the_Unknown_CVPR_2018_paper.html
AUTHORS: Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, Yi Yang
HIGHLIGHT: In this paper, we propose an approach to exploiting unlabeled tracklets by gradually but steadily improving the discriminative capability of the Convolutional Neural Network (CNN) feature representation via stepwise learning.
- 537, TITLE: Pose-Robust Face Recognition via Deep Residual Equivariant Mapping
http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Pose-Robust_Face_Recognition_CVPR_2018_paper.html
AUTHORS: Kaidi Cao, Yu Rong, Cheng Li, Xiaoou Tang, Chen Change Loy
HIGHLIGHT: In this study, we hypothesize that there is an inherent mapping between frontal and profile faces, and consequently, their discrepancy in the deep representation space can be bridged by an equivariant mapping.

538, TITLE: DecideNet: Counting Varying Density Crowds Through Attention Guided Detection and Density Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_DecideNet_Counting_Varying_CVPR_2018_paper.html
AUTHORS: Jiang Liu, Chenqiang Gao, Deyu Meng, Alexander G. Hauptmann
HIGHLIGHT: To address this issue, a novel end-to-end crowd counting framework, named DecideNet (DETeCtIon and Density Estimation Network) is proposed.

539, TITLE: LSTM Pose Machines
http://openaccess.thecvf.com/content_cvpr_2018/html/Luo_LSTM_Pose_Machines_CVPR_2018_paper.html
AUTHORS: Yue Luo, Jimmy Ren, Zhouxia Wang, Wenxiu Sun, Jinshan Pan, Jianbo Liu, Jiahao Pang, Liang Lin
HIGHLIGHT: In this paper, we proposed a novel recurrent network to tackle these problems.

540, TITLE: Disentangling Features in 3D Face Shapes for Joint Face Reconstruction and Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Disentangling_Features_in_CVPR_2018_paper.html
AUTHORS: Feng Liu, Ronghang Zhu, Dan Zeng, Qijun Zhao, Xiaoming Liu
HIGHLIGHT: This paper proposes an encoder-decoder network to disentangle shape features during 3D face shape reconstruction from single 2D images, such that the tasks of learning discriminative shape features for face recognition and reconstructing accurate 3D face shapes can be done simultaneously.

541, TITLE: Convolutional Sequence to Sequence Model for Human Dynamics
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Convolutional_Sequence_to_CVPR_2018_paper.html
AUTHORS: Chen Li, Zhen Zhang, Wee Sun Lee, Gim Hee Lee
HIGHLIGHT: Challenges in modeling human motion include high dimensional prediction as well as extremely complicated dynamics. We present a novel approach to human motion modeling based on convolutional neural networks (CNN).

542, TITLE: Gesture Recognition: Focus on the Hands
http://openaccess.thecvf.com/content_cvpr_2018/html/Narayana_Gesture_Recognition_Focus_CVPR_2018_paper.html
AUTHORS: Pradyumna Narayana, Ross Beveridge, Bruce A. Draper
HIGHLIGHT: We show that when spatial channels are focused on the hands, gesture recognition improves significantly, particularly when the channels are fused using a sparse network.

543, TITLE: Crowd Counting via Adversarial Cross-Scale Consistency Pursuit
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Crowd_Counting_via_CVPR_2018_paper.html
AUTHORS: Zan Shen, Yi Xu, Bingbing Ni, Minsi Wang, Jianguo Hu, Xiaokang Yang
HIGHLIGHT: To explicitly address these issues, we propose a novel crowd counting (density estimation) framework called Adversarial Cross-Scale Consistency Pursuit (ACSCP).

544, TITLE: 3D Human Pose Estimation in the Wild by Adversarial Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_3D_Human_Pose_CVPR_2018_paper.html
AUTHORS: Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, Xiaogang Wang
HIGHLIGHT: In this paper, we propose an adversarial learning framework, which distills the 3D human pose structures learned from the fully annotated dataset to in-the-wild images with only 2D pose annotations.

545, TITLE: CosFace: Large Margin Cosine Loss for Deep Face Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_CosFace_Large_Margin_CVPR_2018_paper.html
AUTHORS: Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, Wei Liu
HIGHLIGHT: In this paper, we propose a novel loss function, namely large margin cosine loss (LMCL), to realize this idea from a different perspective.

546, TITLE: Encoding Crowd Interaction With Deep Neural Network for Pedestrian Trajectory Prediction
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Encoding_Crowd_Interaction_CVPR_2018_paper.html
AUTHORS: Yanyu Xu, Zhixin Piao, Shenghua Gao
HIGHLIGHT: In this paper, we tackle the problem within a deep learning framework by considering motion information of each pedestrian and its interaction with the crowd.

547, TITLE: Mean-Variance Loss for Deep Age Estimation From a Face
http://openaccess.thecvf.com/content_cvpr_2018/html/Pan_Mean-Variance_Loss_for_CVPR_2018_paper.html
AUTHORS: Hongyu Pan, Hu Han, Shiguang Shan, Xilin Chen
HIGHLIGHT: In this paper, we propose a new loss function, called mean-variance loss, for robust age estimation via distribution learning.

- 548, TITLE: Probabilistic Joint Face-Skull Modelling for Facial Reconstruction
http://openaccess.thecvf.com/content_cvpr_2018/html/Madsen_Probabilistic_Joint_Face-Skull_CVPR_2018_paper.html
AUTHORS: Dennis Madsen, Marcel Lüthi, Andreas Schneider, Thomas Vetter
HIGHLIGHT: We present a novel method for co-registration of two independent statistical shape models. Because it is very difficult to obtain the distribution directly from MRI or CT data, we create a dataset of artificial face-skull pairs.
- 549, TITLE: Learning Latent Super-Events to Detect Multiple Activities in Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Piergiovanni_Learning_Latent_Super-Events_CVPR_2018_paper.html
AUTHORS: AJ Piergiovanni, Michael S. Ryoo
HIGHLIGHT: In this paper, we introduce the concept of learning latent super-events from activity videos, and present how it benefits activity detection in continuous videos.
- 550, TITLE: Temporal Hallucinating for Action Recognition With Few Still Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Temporal_Hallucinating_for_CVPR_2018_paper.html
AUTHORS: Yali Wang, Lei Zhou, Yu Qiao
HIGHLIGHT: To mimic this capacity, we propose a novel Hybrid Video Memory (HVM) machine, which can hallucinate temporal features of still images from video memory, in order to boost action recognition with few still images.
- 551, TITLE: Deep Progressive Reinforcement Learning for Skeleton-Based Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Tang_Deep_Progressive_Reinforcement_CVPR_2018_paper.html
AUTHORS: Yansong Tang, Yi Tian, Jiwen Lu, Peiyang Li, Jie Zhou
HIGHLIGHT: In this paper, we propose a deep progressive reinforcement learning (DPRL) method for action recognition in skeleton-based videos, which aims to distill the most informative frames and discard ambiguous frames in sequences for recognizing actions.
- 552, TITLE: Gaze Prediction in Dynamic 360° Immersive Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Gaze_Prediction_in_CVPR_2018_paper.html
AUTHORS: Yanyu Xu, Yanbing Dong, Junru Wu, Zhengzhong Sun, Zhiru Shi, Jingyi Yu, Shenghua Gao
HIGHLIGHT: Considering that the saliency measured at different scales is different, we propose to compute saliency maps at different spatial scales: the sub-image patch centered at current gaze point, the sub-image corresponding to the Field of View (FoV), and the panorama image.
- 553, TITLE: When Will You Do What? - Anticipating Temporal Occurrences of Activities
http://openaccess.thecvf.com/content_cvpr_2018/html/Abu_Farha_When_Will_You_CVPR_2018_paper.html
AUTHORS: Yazan Abu Farha, Alexander Richard, Juergen Gall
HIGHLIGHT: In this paper, we propose two methods to predict a considerably large amount of future actions and their durations.
- 554, TITLE: Fusing Crowd Density Maps and Visual Object Trackers for People Tracking in Crowd Scenes
http://openaccess.thecvf.com/content_cvpr_2018/html/Ren_Fusing_Crowd_Density_CVPR_2018_paper.html
AUTHORS: Weihong Ren, Di Kang, Yandong Tang, Antoni B. Chan
HIGHLIGHT: To train the fusion CNN, we propose a two-stage strategy to gradually optimize the parameters.
- 555, TITLE: Dual Attention Matching Network for Context-Aware Feature Sequence Based Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Si_Dual_Attention_Matching_CVPR_2018_paper.html
AUTHORS: Jianlou Si, Honggang Zhang, Chun-Guang Li, Jason Kuen, Xiangfei Kong, Alex C. Kot, Gang Wang
HIGHLIGHT: In this paper, we propose a novel end-to-end trainable framework, called Dual Attention Matching network (DuATM), to learn context-aware feature sequences and perform attentive sequence comparison simultaneously.
- 556, TITLE: Easy Identification From Better Constraints: Multi-Shot Person Re-Identification From Reference Constraints
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Easy_Identification_From_CVPR_2018_paper.html
AUTHORS: Jiahuan Zhou, Bing Su, Ying Wu
HIGHLIGHT: In this paper, we propose a novel type of similarity constraint.
- 557, TITLE: Crowd Counting With Deep Negative Correlation Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Shi_Crowd_Counting_With_CVPR_2018_paper.html
AUTHORS: Zenglin Shi, Le Zhang, Yun Liu, Xiaofeng Cao, Yangdong Ye, Ming-Ming Cheng, Guoyan Zheng
HIGHLIGHT: Here we propose a new learning strategy to produce generalizable features by way of deep negative correlation learning (NCL).

- 558, TITLE: Human Appearance Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Zanfir_Human_Appearance_Transfer_CVPR_2018_paper.html
AUTHORS: Mihai Zanfir, Alin-Ionut Popa, Andrei Zanfir, Cristian Sminchisescu
HIGHLIGHT: We propose an automatic person-to-person appearance transfer model based on explicit parametric 3d human representations and learned, constrained deep translation network architectures for photographic image synthesis.
- 559, TITLE: Domain Generalization With Adversarial Feature Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Domain_Generalization_With_CVPR_2018_paper.html
AUTHORS: Haoliang Li, Sinno Jialin Pan, Shiqi Wang, Alex C. Kot
HIGHLIGHT: In this paper, we tackle the problem of domain generalization: how to learn a generalized feature representation for an “unseen” target domain by taking the advantage of multiple seen source-domain data.
- 560, TITLE: Pyramid Stereo Matching Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Chang_Pyramid_Stereo_Matching_CVPR_2018_paper.html
AUTHORS: Jia-Ren Chang, Yong-Sheng Chen
HIGHLIGHT: To tackle this problem, we propose PSMNet, a pyramid stereo matching network consisting of two main modules: spatial pyramid pooling and 3D CNN.
- 561, TITLE: Event-Based Vision Meets Deep Learning on Steering Prediction for Self-Driving Cars
http://openaccess.thecvf.com/content_cvpr_2018/html/Maqueda_Event-Based_Vision_Meets_CVPR_2018_paper.html
AUTHORS: Ana I. Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, Davide Scaramuzza
HIGHLIGHT: This paper presents a deep neural network approach that unlocks the potential of event cameras on a challenging motion-estimation task: prediction of a vehicle’s steering angle.
- 562, TITLE: Learning Answer Embeddings for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Learning_Answer_Embeddings_CVPR_2018_paper.html
AUTHORS: Hexiang Hu, Wei-Lun Chao, Fei Sha
HIGHLIGHT: We propose a novel probabilistic model for visual question answering (Visual QA).
- 563, TITLE: Good View Hunting: Learning Photo Composition From Dense View Pairs
http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Good_View_Hunting_CVPR_2018_paper.html
AUTHORS: Zijun Wei, Jianming Zhang, Xiaohui Shen, Zhe Lin, Radomir Mech, Minh Hoai, Dimitris Samaras
HIGHLIGHT: In this work, we present the first large scale Comparative Photo Composition dataset, which contains over one million comparative view pairs annotated using a cost-effective crowdsourcing workflow.
- 564, TITLE: CleanNet: Transfer Learning for Scalable Image Classifier Training With Label Noise
http://openaccess.thecvf.com/content_cvpr_2018/html/Lee_CleanNet_Transfer_Learning_CVPR_2018_paper.html
AUTHORS: Kuang-Huei Lee, Xiaodong He, Lei Zhang, Linjun Yang
HIGHLIGHT: In this paper, we study the problem of learning image classification models with label noise.
- 565, TITLE: Independently Recurrent Neural Network (IndRNN): Building a Longer and Deeper RNN
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Independently_Recurrent_Neural_CVPR_2018_paper.html
AUTHORS: Shuai Li, Wanqing Li, Chris Cook, Ce Zhu, Yanbo Gao
HIGHLIGHT: To address these problems, a new type of RNN, referred to as independently recurrent neural network (IndRNN), is proposed in this paper, where neurons in the same layer are independent of each other and they are connected across layers.
- 566, TITLE: Mix and Match Networks: Encoder-Decoder Alignment for Zero-Pair Image Translation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Mix_and_Match_CVPR_2018_paper.html
AUTHORS: Yaxing Wang, Joost van de Weijer, Luis Herranz
HIGHLIGHT: We address the problem of image translation between domains or modalities for which no direct paired data is available (i.e. zero-pair translation).
- 567, TITLE: Structured Uncertainty Prediction Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Dorta_Structured_Uncertainty_Prediction_CVPR_2018_paper.html
AUTHORS: Garoe Dorta, Sara Vicente, Lourdes Agapito, Neill D. F. Campbell, Ivor Simpson
HIGHLIGHT: Our novel model learns to predict a full Gaussian covariance matrix for each reconstruction, which permits efficient sampling and likelihood evaluation.

- 568, TITLE: Between-Class Learning for Image Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Tokozume_Between-Class_Learning_for_CVPR_2018_paper.html
AUTHORS: Yuji Tokozume, Yoshitaka Ushiku, Tatsuya Harada
HIGHLIGHT: In this paper, we propose a novel learning method for image classification called Between-Class learning (BC learning).
- 569, TITLE: Adversarial Feature Augmentation for Unsupervised Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Volpi_Adversarial_Feature_Augmentation_CVPR_2018_paper.html
AUTHORS: Riccardo Volpi, Pietro Morerio, Silvio Savarese, Vittorio Murino
HIGHLIGHT: In this work, we extend this framework by (i) forcing the learned feature extractor to be domain-invariant, and (ii) training it through data augmentation in the feature space, namely performing feature augmentation.
- 570, TITLE: Generative Image Inpainting With Contextual Attention
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Generative_Image_Inpainting_CVPR_2018_paper.html
AUTHORS: Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, Thomas S. Huang
HIGHLIGHT: Motivated by these observations, we propose a new deep generative model-based approach which can not only synthesize novel image structures but also explicitly utilize surrounding image features as references during network training to make better predictions.
- 571, TITLE: CSGNet: Neural Shape Parser for Constructive Solid Geometry
http://openaccess.thecvf.com/content_cvpr_2018/html/Sharma_CSGNet_Neural_Shape_CVPR_2018_paper.html
AUTHORS: Gopal Sharma, Rishabh Goyal, Difan Liu, Evangelos Kalogerakis, Subhansu Maji
HIGHLIGHT: We present a neural architecture that takes as input a 2D or 3D shape and outputs a program that generates the shape.
- 572, TITLE: Conditional Image-to-Image Translation
http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_Conditional_Image-to-Image_Translation_CVPR_2018_paper.html
AUTHORS: Jianxin Lin, Yingce Xia, Tao Qin, Zhibo Chen, Tie-Yan Liu
HIGHLIGHT: In this paper, we study a new problem, conditional image-to-image translation, which is to translate an image from the source domain to the target domain conditioned on a given image in the target domain.
- 573, TITLE: Continuous Relaxation of MAP Inference: A Nonconvex Perspective
http://openaccess.thecvf.com/content_cvpr_2018/html/Le-Huu_Continuous_Relaxation_of_CVPR_2018_paper.html
AUTHORS: D. Khuê Lê-Huu, Nikos Paragios
HIGHLIGHT: In this paper, we study a nonconvex continuous relaxation of MAP inference in discrete Markov random fields (MRFs).
- 574, TITLE: Feature Generating Networks for Zero-Shot Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Xian_Feature_Generating_Networks_CVPR_2018_paper.html
AUTHORS: Yongqin Xian, Tobias Lorenz, Bernt Schiele, Zeynep Akata
HIGHLIGHT: To circumvent the need for labeled examples of unseen classes, we propose a novel generative adversarial network (GAN) that synthesizes CNN features conditioned on class-level semantic information, offering a shortcut directly from a semantic descriptor of a class to a class-conditional feature distribution.
- 575, TITLE: Joint Optimization Framework for Learning With Noisy Labels
http://openaccess.thecvf.com/content_cvpr_2018/html/Tanaka_Joint_Optimization_Framework_CVPR_2018_paper.html
AUTHORS: Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa
HIGHLIGHT: To overcome this problem, we propose a joint optimization framework of learning DNN parameters and estimating true labels.
- 576, TITLE: Convolutional Image Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Aneja_Convolutional_Image_Captioning_CVPR_2018_paper.html
AUTHORS: Jyoti Aneja, Aditya Deshpande, Alexander G. Schwing
HIGHLIGHT: Inspired by their success, in this paper, we develop a convolutional image captioning technique.
- 577, TITLE: AON: Towards Arbitrarily-Oriented Text Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_AON_Towards_Arbitrarily-Oriented_CVPR_2018_paper.html
AUTHORS: Zhanzhan Cheng, Yangliu Xu, Fan Bai, Yi Niu, Shiliang Pu, Shuigeng Zhou
HIGHLIGHT: In this paper, we develop the arbitrary orientation network (AON) to directly capture the deep features of irregular texts, which are combined into an attention-based decoder to generate character sequence.

- 578, TITLE: Wrapped Gaussian Process Regression on Riemannian Manifolds
http://openaccess.thecvf.com/content_cvpr_2018/html/Mallasto_Wrapped_Gaussian_Process_CVPR_2018_paper.html
AUTHORS: Anton Mallasto, Aasa Feragen
HIGHLIGHT: We tackle the problem by defining wrapped Gaussian processes (WGPs) on Riemannian manifolds, using the probabilistic setting to generalize GP regression to the context of manifold-valued targets.
- 579, TITLE: Geometry Guided Convolutional Neural Networks for Self-Supervised Video Representation Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Gan_Geometry_Guided_Convolutional_CVPR_2018_paper.html
AUTHORS: Chuang Gan, Boqing Gong, Kun Liu, Hao Su, Leonidas J. Guibas
HIGHLIGHT: In this paper, we instead explore geometry, a grand new type of auxiliary supervision for the self-supervised learning of video representations.
- 580, TITLE: DiverseNet: When One Right Answer Is Not Enough
http://openaccess.thecvf.com/content_cvpr_2018/html/Firman_DiverseNet_When_One_CVPR_2018_paper.html
AUTHORS: Michael Firman, Neill D. F. Campbell, Lourdes Agapito, Gabriel J. Brostow
HIGHLIGHT: We introduce a simple method for training a neural network, which enables diverse structured predictions to be made for each test-time query.
- 581, TITLE: Deep Face Detector Adaptation Without Negative Transfer or Catastrophic Forgetting
http://openaccess.thecvf.com/content_cvpr_2018/html/Jamal_Deep_Face_Detector_CVPR_2018_paper.html
AUTHORS: Muhammad Abdullah Jamal, Haoxiang Li, Boqing Gong
HIGHLIGHT: We propose a novel face detector adaptation approach that works as long as there are representative images of the target domain no matter they are labeled or not and, more importantly, without the need of accessing the training data of the source domain.
- 582, TITLE: Analyzing Filters Toward Efficient ConvNet
http://openaccess.thecvf.com/content_cvpr_2018/html/Kobayashi_Analyzing_Filters_Toward_CVPR_2018_paper.html
AUTHORS: Takumi Kobayashi
HIGHLIGHT: In this paper, in contrast to the activations, we focus on filters which are main components of ConvNets.
- 583, TITLE: Regularizing Deep Networks by Modeling and Predicting Label Structure
http://openaccess.thecvf.com/content_cvpr_2018/html/Mostajabi_Regularizing_Deep_Networks_CVPR_2018_paper.html
AUTHORS: Mohammadreza Mostajabi, Michael Maire, Gregory Shakhnarovich
HIGHLIGHT: We construct custom regularization functions for use in supervised training of deep neural networks.
- 584, TITLE: In-Place Activated BatchNorm for Memory-Optimized Training of DNNs
http://openaccess.thecvf.com/content_cvpr_2018/html/Bulo_In-Place_Activated_BatchNorm_CVPR_2018_paper.html
AUTHORS: Samuel Rota Bulò, Lorenzo Porzi, Peter Kotschieder
HIGHLIGHT: In this work we present In-Place Activated Batch Normalization (InPlace-ABN) -- a novel approach to drastically reduce the training memory footprint of modern deep neural networks in a computationally efficient way.
- 585, TITLE: DVQA: Understanding Data Visualizations via Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Kafle_DVQA_Understanding_Data_CVPR_2018_paper.html
AUTHORS: Kushal Kafle, Brian Price, Scott Cohen, Christopher Kanan
HIGHLIGHT: Here, we present DVQA, a dataset that tests many aspects of bar chart understanding in a question answering framework.
- 586, TITLE: DA-GAN: Instance-Level Image Translation by Deep Attention Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Ma_DA-GAN_Instance-Level_Image_CVPR_2018_paper.html
AUTHORS: Shuang Ma, Jianlong Fu, Chang Wen Chen, Tao Mei
HIGHLIGHT: To address the above issues, we propose a novel framework for instance-level image translation by Deep Attention GAN (DA-GAN).
- 587, TITLE: Unsupervised Learning of Depth and Ego-Motion From Monocular Video Using 3D Geometric Constraints
http://openaccess.thecvf.com/content_cvpr_2018/html/Mahjourian_Unsupervised_Learning_of_CVPR_2018_paper.html
AUTHORS: Reza Mahjourian, Martin Wicke, Anelia Angelova
HIGHLIGHT: We present a novel approach for unsupervised learning of depth and ego-motion from monocular video.
- 588, TITLE: FOTS: Fast Oriented Text Spotting With a Unified Network

- http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_FOTS_Fast_Oriented_CVPR_2018_paper.html
AUTHORS: Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, Junjie Yan
HIGHLIGHT: In this work, we propose a unified end-to-end trainable Fast Oriented Text Spotting (FOTS) network for simultaneous detection and recognition, sharing computation and visual information among the two complementary tasks.
- 589, TITLE: Mobile Video Object Detection With Temporally-Aware Feature Maps
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Mobile_Video_Object_CVPR_2018_paper.html
AUTHORS: Mason Liu, Menglong Zhu
HIGHLIGHT: This paper introduces an online model for object detection in videos with real-time performance on mobile and embedded devices.
- 590, TITLE: Weakly Supervised Phrase Localization With Multi-Scale Anchored Transformer Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Weakly_Supervised_Phrase_CVPR_2018_paper.html
AUTHORS: Fang Zhao, Jianshu Li, Jian Zhao, Jiashi Feng
HIGHLIGHT: In this paper, we propose a novel weakly supervised model, Multi-scale Anchored Transformer Network (MATN), to accurately localize free-form textual phrases with only image-level supervision.
- 591, TITLE: Revisiting Oxford and Paris: Large-Scale Image Retrieval Benchmarking
http://openaccess.thecvf.com/content_cvpr_2018/html/Radenovic_Revisiting_Oxford_and_CVPR_2018_paper.html
AUTHORS: Filip Radenovi?, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Ondrej Chum
HIGHLIGHT: In this paper we address issues with image retrieval benchmarking on standard and popular Oxford 5k and Paris 6k datasets.
- 592, TITLE: Cross-Dataset Adaptation for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Chao_Cross-Dataset_Adaptation_for_CVPR_2018_paper.html
AUTHORS: Wei-Lun Chao, Hexiang Hu, Fei Sha
HIGHLIGHT: We overcome this difficulty by proposing a novel domain adaptation algorithm.
- 593, TITLE: Globally Optimal Inlier Set Maximization for Atlanta Frame Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Joo_Globally_Optimal_Inlier_CVPR_2018_paper.html
AUTHORS: Kyungdon Joo, Tae-Hyun Oh, In So Kweon, Jean-Charles Bazin
HIGHLIGHT: In this work, we describe man-made structures via an appropriate structure assumption, called Atlanta world, which contains a vertical direction (typically the gravity direction) and a set of horizontal directions orthogonal to the vertical direction.
- 594, TITLE: End-to-End Convolutional Semantic Embeddings
http://openaccess.thecvf.com/content_cvpr_2018/html/You_End-to-End_Convolutional_Semantic_CVPR_2018_paper.html
AUTHORS: Quanzeng You, Zhengyou Zhang, Jiebo Luo
HIGHLIGHT: In this work, we apply Convolutional Neural Networks to process both images and sentences.
- 595, TITLE: Referring Image Segmentation via Recurrent Refinement Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Referring_Image_Segmentation_CVPR_2018_paper.html
AUTHORS: Ruiyu Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, Jiaya Jia
HIGHLIGHT: In this paper, we utilize the feature pyramids inherently existing in convolutional neural networks to capture the semantics at different scales.
- 596, TITLE: Two Can Play This Game: Visual Dialog With Discriminative Question Generation and Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Jain_Two_Can_Play_CVPR_2018_paper.html
AUTHORS: Unnat Jain, Svetlana Lazebnik, Alexander G. Schwing
HIGHLIGHT: In this paper, we demonstrate a simple symmetric discriminative baseline, that can be applied to both predicting an answer as well as predicting a question.
- 597, TITLE: Generative Adversarial Learning Towards Fast Weakly Supervised Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Generative_Adversarial_Learning_CVPR_2018_paper.html
AUTHORS: Yunhan Shen, Rongrong Ji, Shengchuan Zhang, Wangmeng Zuo, Yan Wang
HIGHLIGHT: In this paper, we speedup online weakly supervised object detectors by orders of magnitude by proposing a novel generative adversarial learning paradigm.
- 598, TITLE: A Deeper Look at Power Normalizations
http://openaccess.thecvf.com/content_cvpr_2018/html/Koniusz_A_Deeper_Look_CVPR_2018_paper.html

AUTHORS: Piotr Koniusz, Hongguang Zhang, Fatih Porikli
HIGHLIGHT: In this paper, we reconsider these operators in the deep learning setup by introducing a novel layer that implements PN for non-linear pooling of feature maps.

599, TITLE: Dimensionality's Blessing: Clustering Images by Underlying Distribution
http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_Dimensionalitys_Blessing_Clustering_CVPR_2018_paper.html
AUTHORS: Wen-Yan Lin, Siying Liu, Jian-Huang Lai, Yasuyuki Matsushita
HIGHLIGHT: We use this to develop distribution-clustering, an elegant algorithm for grouping of data points by their (unknown) underlying distribution.

600, TITLE: Eliminating Background-Bias for Robust Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Tian_Eliminating_Background-Bias_for_CVPR_2018_paper.html
AUTHORS: Maoqing Tian, Shuai Yi, Hongsheng Li, Shihua Li, Xuesen Zhang, Jianping Shi, Junjie Yan, Xiaogang Wang
HIGHLIGHT: To solve the background bias problem, we propose a person-region guided pooling deep neural network based on human parsing maps to learn more discriminative person-part features, and propose to augment training data with person images with random background.

601, TITLE: Learning to Evaluate Image Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Cui_Learning_to_Evaluate_CVPR_2018_paper.html
AUTHORS: Yin Cui, Guandao Yang, Andreas Veit, Xun Huang, Serge Belongie
HIGHLIGHT: To address these two challenges, we propose a novel learning based discriminative evaluation metric that is directly trained to distinguish between human and machine-generated captions.

602, TITLE: Single-Shot Object Detection With Enriched Semantics
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Single-Shot_Object_Detection_CVPR_2018_paper.html
AUTHORS: Zhishuai Zhang, Siyuan Qiao, Cihang Xie, Wei Shen, Bo Wang, Alan L. Yuille
HIGHLIGHT: We propose a novel single shot object detection network named Detection with Enriched Semantics (DES).

603, TITLE: Low-Shot Learning With Imprinted Weights
http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Low-Shot_Learning_With_CVPR_2018_paper.html
AUTHORS: Hang Qi, Matthew Brown, David G. Lowe
HIGHLIGHT: We describe how to add a similar capability to ConvNet classifiers by directly setting the final layer weights from novel training examples during low-shot learning.

604, TITLE: Neural Motifs: Scene Graph Parsing With Global Context
http://openaccess.thecvf.com/content_cvpr_2018/html/Zellers_Neural_Motifs_Scene_CVPR_2018_paper.html
AUTHORS: Rowan Zellers, Mark Yatskar, Sam Thomson, Yejin Choi
HIGHLIGHT: We present new quantitative insights on such repeated structures in the Visual Genome dataset.

605, TITLE: Variational Autoencoders for Deforming 3D Mesh Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Tan_Variational_Autoencoders_for_CVPR_2018_paper.html
AUTHORS: Qingyang Tan, Lin Gao, Yu-Kun Lai, Shihong Xia
HIGHLIGHT: In this paper, we study the problem of analyzing deforming 3D meshes using deep neural networks.

606, TITLE: Fast Monte-Carlo Localization on Aerial Vehicles Using Approximate Continuous Belief Representations
http://openaccess.thecvf.com/content_cvpr_2018/html/Dhawale_Fast_Monte-Carlo_Localization_CVPR_2018_paper.html
AUTHORS: Aditya Dhawale, Kumar Shaurya Shankar, Nathan Michael
HIGHLIGHT: We present a framework to perform fast localization on such platforms enabled by the compressive capabilities of Gaussian Mixture Model representations of point cloud data.

607, TITLE: DeLS-3D: Deep Localization and Segmentation With a 3D Semantic Map
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_DeLS-3D_Deep_Localization_CVPR_2018_paper.html
AUTHORS: Peng Wang, Ruigang Yang, Binbin Cao, Wei Xu, Yuanqing Lin
HIGHLIGHT: In this paper, we propose a unified framework to tackle these two problems simultaneously. In order to validate our approach, we build a dataset with registered 3D point clouds and video camera images.

608, TITLE: LiDAR-Video Driving Dataset: Learning Driving Policies Effectively
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_LiDAR-Video_Driving_Dataset_CVPR_2018_paper.html
AUTHORS: Yiping Chen, Jingkan Wang, Jonathan Li, Cewu Lu, Zhipeng Luo, Han Xue, Cheng Wang

HIGHLIGHT: In this paper, we are the first to propose a LiDAR-Video dataset, which provides large-scale high-quality point clouds scanned by a Velodyne laser, videos recorded by a dashboard camera and standard drivers' behaviors.

609, **TITLE:** Logo Synthesis and Manipulation With Clustered Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Sage_Logo_Synthesis_and_CVPR_2018_paper.html

AUTHORS: Alexander Sage, Eirikur Agustsson, Radu Timofte, Luc Van Gool

HIGHLIGHT: In this paper we explore to what extent machine learning can solve the creative task of the designer. For this, we build a dataset -- LLD -- of 600k+ logos crawled from the world wide web.

610, **TITLE:** Egocentric Basketball Motion Planning From a Single First-Person Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Bertasius_Egocentric_Basketball_Motion_CVPR_2018_paper.html

AUTHORS: Gedas Bertasius, Aaron Chan, Jianbo Shi

HIGHLIGHT: We present a model that uses a single first-person image to generate an egocentric basketball motion sequence in the form of a 12D camera configuration trajectory, which encodes a player's 3D location and 3D head orientation throughout the sequence.

611, **TITLE:** Human-Centric Indoor Scene Synthesis Using Stochastic Grammar

http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Human-Centric_Indoor_Scene_CVPR_2018_paper.html

AUTHORS: Siyuan Qi, Yixin Zhu, Siyuan Huang, Chenfanfu Jiang, Song-Chun Zhu

HIGHLIGHT: We present a human-centric method to sample and synthesize 3D room layouts and 2D images thereof, for the purpose of obtaining large-scale 2D/3D image data with the perfect per-pixel ground truth.

612, **TITLE:** Rotation-Sensitive Regression for Oriented Scene Text Detection

http://openaccess.thecvf.com/content_cvpr_2018/html/Liao_Rotation-Sensitive_Regression_for_CVPR_2018_paper.html

AUTHORS: Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, Xiang Bai

HIGHLIGHT: To address this issue, we propose to perform classification and regression on features of different characteristics, extracted by two network branches of different designs.

613, **TITLE:** Separating Self-Expression and Visual Content in Hashtag Supervision

http://openaccess.thecvf.com/content_cvpr_2018/html/Veit_Separating_Self-Expression_and_CVPR_2018_paper.html

AUTHORS: Andreas Veit, Maximilian Nickel, Serge Belongie, Laurens van der Maaten

HIGHLIGHT: This paper presents an approach that extends upon modeling simple image-label pairs with a joint model of images, hashtags, and users.

614, **TITLE:** Distort-and-Recover: Color Enhancement Using Deep Reinforcement Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Park_Distort-and-Recover_Color_Enhancement_CVPR_2018_paper.html

AUTHORS: Jongchan Park, Joon-Young Lee, Donggeun Yoo, In So Kweon

HIGHLIGHT: In this paper, we present a deep reinforcement learning (DRL) based method for color enhancement to explicitly model the step-wise nature of human retouching process.

615, **TITLE:** Im2Flow: Motion Hallucination From Static Images for Action Recognition

http://openaccess.thecvf.com/content_cvpr_2018/html/Gao_Im2Flow_Motion_Hallucination_CVPR_2018_paper.html

AUTHORS: Ruohan Gao, Bo Xiong, Kristen Grauman

HIGHLIGHT: We propose an approach that hallucinates the unobserved future motion implied by a single snapshot to help static-image action recognition.

616, **TITLE:** Finding "It": Weakly-Supervised Reference-Aware Visual Grounding in Instructional Videos

http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Finding_It_Weakly-Supervised_CVPR_2018_paper.html

AUTHORS: De-An Huang, Shyamal Buch, Lucio Dery, Animesh Garg, Li Fei-Fei, Juan Carlos Niebles

HIGHLIGHT: In this work, we propose to tackle this new task with a weakly-supervised framework for reference-aware visual grounding in instructional videos, where only the temporal alignment between the transcription and the video segment are available for supervision.

617, **TITLE:** Actor and Action Video Segmentation From a Sentence

http://openaccess.thecvf.com/content_cvpr_2018/html/Gavrilyuk_Actor_and_Action_CVPR_2018_paper.html

AUTHORS: Kirill Gavrilyuk, Amir Ghodrati, Zhenyang Li, Cees G. M. Snoek

HIGHLIGHT: We propose a fully-convolutional model for pixel-level actor and action segmentation using an encoder-decoder architecture optimized for video.

618, **TITLE:** Egocentric Activity Recognition on a Budget

http://openaccess.thecvf.com/content_cvpr_2018/html/Possas_Egocentric_Activity_Recognition_CVPR_2018_paper.html
AUTHORS: Rafael Possas, Sheila Pinto Caceres, Fabio Ramos
HIGHLIGHT: We develop a Reinforcement Learning model-free method to learn energy-aware policies that maximize the use of low-energy cost predictors while keeping competitive accuracy levels.

619, TITLE: CNN in MRF: Video Object Segmentation via Inference in a CNN-Based Higher-Order Spatio-Temporal MRF
http://openaccess.thecvf.com/content_cvpr_2018/html/Bao_CNN_in_MRF_CVPR_2018_paper.html
AUTHORS: Linchao Bao, Baoyuan Wu, Wei Liu
HIGHLIGHT: We propose a novel spatio-temporal Markov Random Field (MRF) model defined over pixels to handle this problem.

620, TITLE: Action Sets: Weakly Supervised Action Segmentation Without Ordering Constraints
http://openaccess.thecvf.com/content_cvpr_2018/html/Richard_Action_Sets_Weakly_CVPR_2018_paper.html
AUTHORS: Alexander Richard, Hilde Kuehne, Juergen Gall
HIGHLIGHT: We introduce a system that automatically learns to temporally segment and label actions in a video, where the only supervision that is used are action sets.

621, TITLE: Low-Latency Video Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Low-Latency_Video_Semantic_CVPR_2018_paper.html
AUTHORS: Yule Li, Jianping Shi, Dahua Lin
HIGHLIGHT: To tackle this combined challenge, we develop a framework for video semantic segmentation, which incorporates two novel components:(1) a feature propagation module that adaptively fuses features over time via spatially variant convolution, thus reducing the cost of per-frame computation; and (2) an adaptive scheduler that dynamically allocate computation based on accuracy prediction.

622, TITLE: Fine-Grained Video Captioning for Sports Narrative
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Fine-Grained_Video_Captioning_CVPR_2018_paper.html
AUTHORS: Huanyu Yu, Shuo Cheng, Bingbing Ni, Minsi Wang, Jian Zhang, Xiaokang Yang
HIGHLIGHT: To this end, this work makes the following contributions.
First, to facilitate this novel research of fine-grained video caption, we collected a novel dataset called Fine-grained Sports Narrative dataset (FSN) that contains 2K sports videos with ground-truth narratives from YouTube.com.

623, TITLE: End-to-End Learning of Motion Representation for Video Understanding
http://openaccess.thecvf.com/content_cvpr_2018/html/Fan_End-to-End_Learning_of_CVPR_2018_paper.html
AUTHORS: Lijie Fan, Wenbing Huang, Chuang Gan, Stefano Ermon, Boqing Gong, Junzhou Huang
HIGHLIGHT: To fill this gap, we propose TVNet, a novel end-to-end trainable neural network, to learn optical-flow-like features from data.

624, TITLE: Compressed Video Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Compressed_Video_Action_CVPR_2018_paper.html
AUTHORS: Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R. Manmatha, Alexander J. Smola, Philipp Krähenbühl
HIGHLIGHT: We propose novel techniques to use them effectively.

625, TITLE: Features for Multi-Target Multi-Camera Tracking and Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Ristani_Features_for_Multi-Target_CVPR_2018_paper.html
AUTHORS: Ergys Ristani, Carlo Tomasi
HIGHLIGHT: Our contributions include an adaptive weighted triplet loss for training and a new technique for hard-identity mining.

626, TITLE: AVA: A Video Dataset of Spatio-Temporally Localized Atomic Visual Actions
http://openaccess.thecvf.com/content_cvpr_2018/html/Gu_AVA_A_Video_CVPR_2018_paper.html
AUTHORS: Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, Jitendra Malik
HIGHLIGHT: This paper introduces a video dataset of spatio-temporally localized Atomic Visual Actions (AVA).

627, TITLE: Who's Better? Who's Best? Pairwise Deep Ranking for Skill Determination
http://openaccess.thecvf.com/content_cvpr_2018/html/Doughty_Whos_Better_Whos_CVPR_2018_paper.html
AUTHORS: Hazel Doughty, Dima Damen, Walterio Mayol-Cuevas
HIGHLIGHT: This paper presents a method for assessing skill from video, applicable to a variety of tasks, ranging from surgery to drawing and rolling pizza dough.

- 628, TITLE: MX-LSTM: Mixing Tracklets and Vislets to Jointly Forecast Trajectories and Head Poses
http://openaccess.thecvf.com/content_cvpr_2018/html/Hasan_MX-LSTM_Mixing_Tracklets_CVPR_2018_paper.html
AUTHORS: Irtiza Hasan, Francesco Setti, Theodore Tsesmelis, Alessio Del Bue, Fabio Galasso, Marco Cristani
HIGHLIGHT: This paper shows that adding vislets, that is, short sequences of head pose estimations, allows to increase significantly the trajectory forecasting performance.
- 629, TITLE: Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Anderson_Bottom-Up_and_Top-Down_CVPR_2018_paper.html
AUTHORS: Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, Lei Zhang
HIGHLIGHT: In this work, we propose a combined bottom-up and top-down attention mechanism that enables attention to be calculated at the level of objects and other salient image regions.
- 630, TITLE: Improved Fusion of Visual and Language Representations by Dense Symmetric Co-Attention for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Nguyen_Improved_Fusion_of_CVPR_2018_paper.html
AUTHORS: Duy-Kien Nguyen, Takayuki Okatani
HIGHLIGHT: Specifically, we present a simple architecture that is fully symmetric between visual and language representations, in which each question word attends on image regions and each image region attends on question words.
- 631, TITLE: FlipDial: A Generative Model for Two-Way Visual Dialogue
http://openaccess.thecvf.com/content_cvpr_2018/html/Massiceti_FlipDial_A_Generative_CVPR_2018_paper.html
AUTHORS: Daniela Massiceti, N. Siddharth, Puneet K. Dokania, Philip H.S. Torr
HIGHLIGHT: We present FlipDial, a generative model for Visual Dialogue that simultaneously plays the role of both participants in a visually-grounded dialogue.
We are the first to extend this paradigm to full two-way visual dialogue (2VD), where our model is capable of generating both questions and answers in sequence based on a visual input, for which we propose a set of novel evaluation measures and metrics.
- 632, TITLE: Are You Talking to Me? Reasoned Visual Dialog Generation Through Adversarial Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Are_You_Talking_CVPR_2018_paper.html
AUTHORS: Qi Wu, Peng Wang, Chunhua Shen, Ian Reid, Anton van den Hengel
HIGHLIGHT: We present a novel approach that combines Reinforcement Learning and Generative Adversarial Networks (GANs) to generate more human-like responses to questions.
- 633, TITLE: Visual Question Generation as Dual Task of Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Visual_Question_Generation_CVPR_2018_paper.html
AUTHORS: Yikang Li, Nan Duan, Bolei Zhou, Xiao Chu, Wanli Ouyang, Xiaogang Wang, Ming Zhou
HIGHLIGHT: In this paper, we propose an end-to-end unified model, the Invertible Question Answering Network (iQAN), to introduce question generation as a dual task of question answering to improve the VQA performance.
- 634, TITLE: Unsupervised Textual Grounding: Linking Words to Image Concepts
http://openaccess.thecvf.com/content_cvpr_2018/html/Yeh_Unsupervised_Textual_Grounding_CVPR_2018_paper.html
AUTHORS: Raymond A. Yeh, Minh N. Do, Alexander G. Schwing
HIGHLIGHT: We demonstrate our approach on the ReferIt Game dataset and the Flickr30k data, outperforming baselines by 7.98% and 6.96% respectively.
- 635, TITLE: Focal Visual-Text Attention for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Liang_Focal_Visual-Text_Attention_CVPR_2018_paper.html
AUTHORS: Junwei Liang, Lu Jiang, Liangliang Cao, Li-Jia Li, Alexander G. Hauptmann
HIGHLIGHT: In this paper, we describe a novel neural network called Focal Visual-Text Attention network (FVTA) for collective reasoning in visual question answering, where both visual and text sequence information such as images and text metadata are presented.
- 636, TITLE: SeGAN: Segmenting and Generating the Invisible
http://openaccess.thecvf.com/content_cvpr_2018/html/Ehsani_SeGAN_Segmenting_and_CVPR_2018_paper.html
AUTHORS: Kiana Ehsani, Roozbeh Mottaghi, Ali Farhadi
HIGHLIGHT: In this paper, we study the challenging problem of completing the appearance of occluded objects.
- 637, TITLE: Cascade R-CNN: Delving Into High Quality Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Cai_Cascade_R-CNN_Delving_CVPR_2018_paper.html
AUTHORS: Zhaowei Cai, Nuno Vasconcelos

HIGHLIGHT: Cascade R-CNN: Delving Into High Quality Object Detection

638, TITLE: Learning Semantic Concepts and Order for Image and Sentence Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Learning_Semantic_Concepts_CVPR_2018_paper.html

AUTHORS: Yan Huang, Qi Wu, Chunfeng Song, Liang Wang

HIGHLIGHT: In this work, we propose a semantic-enhanced image and sentence matching model, which can improve the image representation by learning semantic concepts and then organizing them in a correct semantic order.

639, TITLE: Functional Map of the World
http://openaccess.thecvf.com/content_cvpr_2018/html/Christie_Functional_Map_of_CVPR_2018_paper.html

AUTHORS: Gordon Christie, Neil Fendley, James Wilson, Ryan Mukherjee

HIGHLIGHT: We present an analysis of the dataset along with baseline approaches that reason about metadata and temporal views.

We present a new dataset, Functional Map of the World (fMoW), which aims to inspire the development of machine learning models capable of predicting the functional purpose of buildings and land use from temporal sequences of satellite images and a rich set of metadata features.

640, TITLE: MegDet: A Large Mini-Batch Object Detector
http://openaccess.thecvf.com/content_cvpr_2018/html/Peng_MegDet_A_Large_CVPR_2018_paper.html

AUTHORS: Chao Peng, Tete Xiao, Zeming Li, Yuning Jiang, Xiangyu Zhang, Kai Jia, Gang Yu, Jian Sun

HIGHLIGHT: In this paper, we propose a Large Mini-Batch Object Detector (MegDet) to enable the training with a large mini-batch size up to 256, so that we can effectively utilize at most 128 GPUs to significantly shorten the training time.

641, TITLE: Learning Globally Optimized Object Detector via Policy Gradient
http://openaccess.thecvf.com/content_cvpr_2018/html/Rao_Learning_Globally_Optimized_CVPR_2018_paper.html

AUTHORS: Yongming Rao, Dahua Lin, Jiwen Lu, Jie Zhou

HIGHLIGHT: In this paper, we propose a simple yet effective method to learn globally optimized detector for object detection, which is a simple modification to the standard cross-entropy gradient inspired by the REINFORCE algorithm.

642, TITLE: Photographic Text-to-Image Synthesis With a Hierarchically-Nested Adversarial Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Photographic_Text-to-Image_Synthesis_CVPR_2018_paper.html

AUTHORS: Zizhao Zhang, Yuanpu Xie, Lin Yang

HIGHLIGHT: This paper presents a novel method to deal with the challenging task of generating photographic images conditioned on semantic image descriptions.

643, TITLE: Illuminant Spectra-Based Source Separation Using Flash Photography
http://openaccess.thecvf.com/content_cvpr_2018/html/Hui_Illuminant_Spectra-Based_Source_CVPR_2018_paper.html

AUTHORS: Zhuo Hui, Kalyan Sunkavalli, Sunil Hadap, Aswin C. Sankaranarayanan

HIGHLIGHT: In this work, we leverage a flash/no-flash image pair to analyze and edit scene illuminants based on their spectral differences.

644, TITLE: Trapping Light for Time of Flight
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Trapping_Light_for_CVPR_2018_paper.html

AUTHORS: Ruilin Xu, Mohit Gupta, Shree K. Nayar

HIGHLIGHT: We propose a novel imaging method for near-complete, surround, 3D reconstruction of geometrically complex objects, in a single shot.

645, TITLE: The Perception-Distortion Tradeoff
http://openaccess.thecvf.com/content_cvpr_2018/html/Blau_The_Perception-Distortion_Tradeoff_CVPR_2018_paper.html

AUTHORS: Yochai Blau, Tomer Michaeli

HIGHLIGHT: In this paper, we prove mathematically that distortion and perceptual quality are at odds with each other.

646, TITLE: Label Denoising Adversarial Network (LDAN) for Inverse Lighting of Faces
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Label_Denoising_Adversarial_CVPR_2018_paper.html

AUTHORS: Hao Zhou, Jin Sun, Yaser Yacoob, David W. Jacobs

HIGHLIGHT: We propose to train a deep Convolutional Neural Network (CNN) to regress lighting parameters from a single face image.

647, TITLE: Optimal Structured Light à La Carte
http://openaccess.thecvf.com/content_cvpr_2018/html/Mirdehghan_Optimal_Structured_Light_CVPR_2018_paper.html

- AUTHORS: Parsa Mirdehghan, Wenzheng Chen, Kiriakos N. Kutulakos
HIGHLIGHT: We consider the problem of automatically generating sequences of structured-light patterns for active stereo triangulation of a static scene.
- 648, TITLE: Tracking Multiple Objects Outside the Line of Sight Using Speckle Imaging
http://openaccess.thecvf.com/content_cvpr_2018/html/Smith_Tracking_Multiple_Objects_CVPR_2018_paper.html
AUTHORS: Brandon M. Smith, Matthew O'Toole, Mohit Gupta
HIGHLIGHT: This paper presents techniques for tracking non-line-of-sight (NLOS) objects using speckle imaging.
- 649, TITLE: Inferring Light Fields From Shadows
http://openaccess.thecvf.com/content_cvpr_2018/html/Baradad_Inferring_Light_Fields_CVPR_2018_paper.html
AUTHORS: Manel Baradad, Vickie Ye, Adam B. Yedidia, Frédo Durand, William T. Freeman, Gregory W. Wornell, Antonio Torralba
HIGHLIGHT: We present a method for inferring a 4D light field of a hidden scene from 2D shadows cast by a known occluder on a diffuse wall.
- 650, TITLE: Modifying Non-Local Variations Across Multiple Views
http://openaccess.thecvf.com/content_cvpr_2018/html/TLusty_Modifying_Non-Local_Variations_CVPR_2018_paper.html
AUTHORS: Tal Tlusty, Tomer Michaeli, Tali Dekel, Lih Zelnik-Manor
HIGHLIGHT: We present an algorithm for modifying small non-local variations between repeating structures and patterns in multiple images of the same scene.
- 651, TITLE: Robust Video Content Alignment and Compensation for Rain Removal in a CNN Framework
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Robust_Video_Content_CVPR_2018_paper.html
AUTHORS: Jie Chen, Cheen-Hau Tan, Junhui Hou, Lap-Pui Chau, He Li
HIGHLIGHT: We propose a novel derain algorithm, which applies superpixel (SP) segmentation to decompose the scene into depth consistent units.
- 652, TITLE: SfSNet: Learning Shape, Reflectance and Illuminance of Faces 'in the Wild'
http://openaccess.thecvf.com/content_cvpr_2018/html/Sengupta_SfSNet_Learning_Shape_CVPR_2018_paper.html
AUTHORS: Soumyadip Sengupta, Angjoo Kanazawa, Carlos D. Castillo, David W. Jacobs
HIGHLIGHT: We present SfSNet, an end-to-end learning framework for producing an accurate decomposition of an unconstrained human face image into shape, reflectance and illuminance.
- 653, TITLE: Deep Photo Enhancer: Unpaired Learning for Image Enhancement From Photographs With GANs
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Deep_Photo_Enhancer_CVPR_2018_paper.html
AUTHORS: Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, Yung-Yu Chuang
HIGHLIGHT: This paper proposes an unpaired learning method for image enhancement.
- 654, TITLE: LIME: Live Intrinsic Material Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Meka_LIME_Live_Intrinsic_CVPR_2018_paper.html
AUTHORS: Abhimitra Meka, Maxim Maximov, Michael Zollhöfer, Avishek Chatterjee, Hans-Peter Seidel, Christian Richardt, Christian Theobalt
HIGHLIGHT: We present the first end-to-end approach for real-time material estimation for general object shapes with uniform material that only requires a single color image as input.
- 655, TITLE: Learning to Detect Features in Texture Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Learning_to_Detect_CVPR_2018_paper.html
AUTHORS: Linguang Zhang, Szymon Rusinkiewicz
HIGHLIGHT: We propose an effective and scalable method for learning feature detectors for textures, which combines an existing "ranking" loss with an efficient fully-convolutional architecture as well as a new training-loss term that maximizes the "peakedness" of the response map.
- 656, TITLE: Learning to Extract a Video Sequence From a Single Motion-Blurred Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Jin_Learning_to_Extract_CVPR_2018_paper.html
AUTHORS: Meiguang Jin, Givi Meishvili, Paolo Favaro
HIGHLIGHT: We present a method to extract a video sequence from a single motion-blurred image.
- 657, TITLE: Lose the Views: Limited Angle CT Reconstruction via Implicit Sinogram Completion
http://openaccess.thecvf.com/content_cvpr_2018/html/Anirudh_Lose_the_Views_CVPR_2018_paper.html

AUTHORS: Rushil Anirudh, Hyojin Kim, Jayaraman J. Thiagarajan, K. Aditya Mohan, Kyle Champley, Timo Bremer
HIGHLIGHT: In this paper, we propose to address this problem using CTNet -- a system of 1D and 2D convolutional neural networks, that operates directly on a limited angle sinogram to predict the reconstruction.

658, TITLE: A Common Framework for Interactive Texture Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Men_A_Common_Framework_CVPR_2018_paper.html
AUTHORS: Yifang Men, Zhouhui Lian, Yingmin Tang, Jianguo Xiao
HIGHLIGHT: In this paper, we present a general-purpose solution to interactive texture transfer problems that better preserves both local structure and visual richness.

659, TITLE: AMNet: Memorability Estimation With Attention
http://openaccess.thecvf.com/content_cvpr_2018/html/Fajtl_AMNet_Memorability_Estimation_CVPR_2018_paper.html
AUTHORS: Jiri Fajtl, Vasileios Argyriou, Dorothy Monekosso, Paolo Remagnino
HIGHLIGHT: In this paper we present the design and evaluation of an end to end trainable, deep neural network with a visual attention mechanism for memorability estimation in still images.

660, TITLE: Blind Predicting Similar Quality Map for Image Quality Assessment
http://openaccess.thecvf.com/content_cvpr_2018/html/Pan_Blind_Predicting_Similar_CVPR_2018_paper.html
AUTHORS: Da Pan, Ping Shi, Ming Hou, Zefeng Ying, Sizhe Fu, Yuan Zhang
HIGHLIGHT: In this paper, we propose a simple and efficient BIQA model based on a novel framework which consists of a fully convolutional neural network (FCNN) and a pooling network to solve this problem.

661, TITLE: Deep End-to-End Time-of-Flight Imaging
http://openaccess.thecvf.com/content_cvpr_2018/html/Su_Deep_End-to-End_Time-of-Flight_CVPR_2018_paper.html
AUTHORS: Shuo Chen Su, Felix Heide, Gordon Wetzstein, Wolfgang Heidrich
HIGHLIGHT: We present an end-to-end image processing framework for time-of-flight (ToF) cameras.

662, TITLE: Aperture Supervision for Monocular Depth Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Srinivasan_Aperture_Supervision_for_CVPR_2018_paper.html
AUTHORS: Pratul P. Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, Jonathan T. Barron
HIGHLIGHT: We present a novel method to train machine learning algorithms to estimate scene depths from a single image, by using the information provided by a camera's aperture as supervision.

663, TITLE: Seeing Temporal Modulation of Lights From Standard Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Sakakibara_Seeing_Temporal_Modulation_CVPR_2018_paper.html
AUTHORS: Naoki Sakakibara, Fumihiko Sakaue, Jun Sato
HIGHLIGHT: In this paper, we propose a novel method for measuring the temporal modulation of lights by using off-the-shelf cameras.

664, TITLE: Statistical Tomography of Microscopic Life
http://openaccess.thecvf.com/content_cvpr_2018/html/Levis_Statistical_Tomography_of_CVPR_2018_paper.html
AUTHORS: Aviad Levis, Yoav Y. Schechner, Ronen Talmon
HIGHLIGHT: We apply the method to study of plankton.

665, TITLE: Divide and Conquer for Full-Resolution Light Field Deblurring
http://openaccess.thecvf.com/content_cvpr_2018/html/Mohan_Divide_and_Conquer_CVPR_2018_paper.html
AUTHORS: M. R. Mahesh Mohan, A. N. Rajagopalan
HIGHLIGHT: In this paper, we introduce a new blind motion deblurring strategy for LFs which alleviates these limitations significantly.

666, TITLE: Multispectral Image Intrinsic Decomposition via Subspace Constraint
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Multispectral_Image_Intrinsic_CVPR_2018_paper.html
AUTHORS: Qian Huang, Weixin Zhu, Yang Zhao, Linsen Chen, Yao Wang, Tao Yue, Xun Cao
HIGHLIGHT: In this paper, a new Multispectral Image Intrinsic Decomposition model (MIID) is presented to decompose the shading and reflectance from a single multispectral image.

667, TITLE: Improving Color Reproduction Accuracy on Cameras
http://openaccess.thecvf.com/content_cvpr_2018/html/Karaimer_Improving_Color_Reproduction_CVPR_2018_paper.html
AUTHORS: Hakki Can Karaimer, Michael S. Brown

HIGHLIGHT: In this paper, we discuss the limitations of the current colorimetric mapping approach and propose two methods that are able to improve color accuracy.

668, **TITLE:** A Closer Look at Spatiotemporal Convolutions for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Tran_A_Closer_Look_CVPR_2018_paper.html
AUTHORS: Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, Manohar Paluri
HIGHLIGHT: In this paper we discuss several forms of spatiotemporal convolutions for video analysis and study their effects on action recognition.

669, **TITLE:** Inferring Shared Attention in Social Scene Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Fan_Inferring_Shared_Attention_CVPR_2018_paper.html
AUTHORS: Lifeng Fan, Yixin Chen, Ping Wei, Wenguan Wang, Song-Chun Zhu
HIGHLIGHT: We propose a spatial-temporal neural network to detect shared attention intervals in videos and predict shared attention locations in frames.

670, **TITLE:** Making Convolutional Networks Recurrent for Visual Sequence Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Making_Convolutional_Networks_CVPR_2018_paper.html
AUTHORS: Xiaodong Yang, Pavlo Molchanov, Jan Kautz
HIGHLIGHT: In this paper, we aim to bridge this gap and present the first large-scale exploration of RNNs for visual sequence learning.

671, **TITLE:** Real-World Anomaly Detection in Surveillance Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Sultani_Real-World_Anomaly_Detection_CVPR_2018_paper.html
AUTHORS: Waqas Sultani, Chen Chen, Mubarak Shah
HIGHLIGHT: In this paper, we propose to learn anomalies by exploiting both normal and anomalous videos.

672, **TITLE:** Viewpoint-Aware Attentive Multi-View Inference for Vehicle Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Viewpoint-Aware_Attentive_Multi-View_CVPR_2018_paper.html
AUTHORS: Yi Zhou, Ling Shao
HIGHLIGHT: In this paper, we propose a Viewpoint-aware Attentive Multi-view Inference (VAMI) model that only requires visual information to solve the multi-view vehicle re-ID problem.

673, **TITLE:** Efficient Video Object Segmentation via Network Modulation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Efficient_Video_Object_CVPR_2018_paper.html
AUTHORS: Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, Aggelos K. Katsaggelos
HIGHLIGHT: We propose a novel approach that uses a single forward pass to adapt the segmentation model to the appearance of a specific object.

674, **TITLE:** Weakly-Supervised Action Segmentation With Iterative Soft Boundary Assignment
http://openaccess.thecvf.com/content_cvpr_2018/html/Ding_Weakly-Supervised_Action_Segmentation_CVPR_2018_paper.html
AUTHORS: Li Ding, Chenliang Xu
HIGHLIGHT: In this work, we address the task of weakly-supervised human action segmentation in long, untrimmed videos.

675, **TITLE:** Depth-Aware Stereo Video Retargeting
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Depth-Aware_Stereo_Video_CVPR_2018_paper.html
AUTHORS: Bing Li, Chia-Wen Lin, Boxin Shi, Tiejun Huang, Wen Gao, C.-C. Jay Kuo
HIGHLIGHT: In this work, we propose a depth-aware stereo video retargeting method by imposing the depth fidelity constraint.

676, **TITLE:** Instance Embedding Transfer to Unsupervised Video Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Instance_Embedding_Transfer_CVPR_2018_paper.html
AUTHORS: Siyang Li, Bryan Seybold, Alexey Vorobyov, Alireza Fathi, Qin Huang, C.-C. Jay Kuo
HIGHLIGHT: We propose a method for unsupervised video object segmentation by transferring the knowledge encapsulated in image-based instance embedding networks.

677, **TITLE:** Future Frame Prediction for Anomaly Detection – A New Baseline
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Future_Frame_Prediction_CVPR_2018_paper.html
AUTHORS: Wen Liu, Weixin Luo, Dongze Lian, Shenghua Gao
HIGHLIGHT: In this paper, we propose to tackle the anomaly detection problem within a video prediction framework.

678, TITLE: Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet?
http://openaccess.thecvf.com/content_cvpr_2018/html/Hara_Can_Spatiotemporal_3D_CVPR_2018_paper.html
AUTHORS: Kensho Hara, Hirokatsu Kataoka, Yutaka Satoh
HIGHLIGHT: The purpose of this study is to determine whether current video datasets have sufficient data for training very deep convolutional neural networks (CNNs) with spatio-temporal three-dimensional (3D) kernels.

679, TITLE: Dynamic Video Segmentation Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Dynamic_Video_Segmentation_CVPR_2018_paper.html
AUTHORS: Yu-Syuan Xu, Tsu-Jui Fu, Hsuan-Kung Yang, Chun-Yi Lee
HIGHLIGHT: In this paper, we present a detailed design of dynamic video segmentation network (DVSNet) for fast and efficient semantic video segmentation.

680, TITLE: Recognize Actions by Disentangling Components of Dynamics
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Recognize_Actions_by_CVPR_2018_paper.html
AUTHORS: Yue Zhao, Yuanjun Xiong, Dahua Lin
HIGHLIGHT: In this paper, we propose a new ConvNet architecture for video representation learning, which can derive disentangled components of dynamics purely from raw video frames, without the need of optical flow estimation.

681, TITLE: Motion-Appearance Co-Memory Networks for Video Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Gao_Motion-Appearance_Co-Memory_Networks_CVPR_2018_paper.html
AUTHORS: Jiyang Gao, Runzhou Ge, Kan Chen, Ram Nevatia
HIGHLIGHT: Based these observations, we propose a motion-appearance co-memory network for video QA.

682, TITLE: Learning to Understand Image Blur
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Learning_to_Understand_CVPR_2018_paper.html
AUTHORS: Shanghang Zhang, Xiaohui Shen, Zhe Lin, Radomír Měch, João P. Costeira, José M. F. Moura
HIGHLIGHT: In this paper, we propose a unified framework to estimate a spatially-varying blur map and understand its desirability in terms of image quality at the same time.
Considering the limitations of existing image blur datasets, we collected a new large-scale dataset with both annotations to facilitate training.

683, TITLE: Dense Decoder Shortcut Connections for Single-Pass Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Bilinski_Dense_Decoder_Shortcut_CVPR_2018_paper.html
AUTHORS: Piotr Bilinski, Victor Prisacariu
HIGHLIGHT: We propose a novel end-to-end trainable, deep, encoder-decoder architecture for single-pass semantic segmentation.

684, TITLE: Generative Adversarial Image Synthesis With Decision Tree Latent Controller
http://openaccess.thecvf.com/content_cvpr_2018/html/Kaneko_Generative_Adversarial_Image_CVPR_2018_paper.html
AUTHORS: Takuhiro Kaneko, Kaoru Hiramatsu, Kunio Kashino
HIGHLIGHT: This paper proposes the decision tree latent controller generative adversarial network (DTLC-GAN), an extension of a GAN that can learn hierarchically interpretable representations without relying on detailed supervision.

685, TITLE: Learning a Discriminative Prior for Blind Image Deblurring
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Learning_a_Discriminative_CVPR_2018_paper.html
AUTHORS: Lerenhan Li, Jinshan Pan, Wei-Sheng Lai, Changxin Gao, Nong Sang, Ming-Hsuan Yang
HIGHLIGHT: We present an effective blind image deblurring method based on a data-driven discriminative prior.

686, TITLE: Frame-Recurrent Video Super-Resolution
http://openaccess.thecvf.com/content_cvpr_2018/html/Sajjadi_Frame-Recurrent_Video_Super-Resolution_CVPR_2018_paper.html
AUTHORS: Mehdi S. M. Sajjadi, Raviteja Vemulapalli, Matthew Brown
HIGHLIGHT: In this work, we propose an end-to-end trainable frame-recurrent video super-resolution framework that uses the previously inferred HR estimate to super-resolve the subsequent frame.

687, TITLE: Discovering Point Lights With Intensity Distance Fields
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Discovering_Point_Lights_CVPR_2018_paper.html
AUTHORS: Edward Zhang, Michael F. Cohen, Brian Curless
HIGHLIGHT: We introduce the light localization problem.

- 688, TITLE: Video Rain Streak Removal by Multiscale Convolutional Sparse Coding
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Video_Rain_Streak_CVPR_2018_paper.html
AUTHORS: Minghan Li, Qi Xie, Qian Zhao, Wei Wei, Shuhang Gu, Jing Tao, Deyu Meng
HIGHLIGHT: In this paper, we raise two intrinsic characteristics specifically possessed by rain streaks.
- 689, TITLE: Stereoscopic Neural Style Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Stereoscopic_Neural_Style_CVPR_2018_paper.html
AUTHORS: Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, Gang Hua
HIGHLIGHT: For a practical real-time solution, we propose the first feed-forward network by jointly training a stylization sub-network and a disparity sub-network, and integrate them in a feature level middle domain.
- 690, TITLE: Multi-Frame Quality Enhancement for Compressed Video
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Multi-Frame_Quality_Enhancement_CVPR_2018_paper.html
AUTHORS: Ren Yang, Mai Xu, Zulin Wang, Tianyi Li
HIGHLIGHT: In this paper, we investigate that heavy quality fluctuation exists across compressed video frames, and thus low quality frames can be enhanced using the neighboring high quality frames, seen as Multi-Frame Quality Enhancement (MFQE).
- 691, TITLE: CNN Based Learning Using Reflection and Retinex Models for Intrinsic Image Decomposition
http://openaccess.thecvf.com/content_cvpr_2018/html/Baslamisli_CNN_Based_Learning_CVPR_2018_paper.html
AUTHORS: Anil S. Baslamisli, Hoang-An Le, Theo Gevers
HIGHLIGHT: In this paper, the aim is to exploit the best of the two worlds.
- 692, TITLE: Image Restoration by Estimating Frequency Distribution of Local Patches
http://openaccess.thecvf.com/content_cvpr_2018/html/Yoo_Image_Restoration_by_CVPR_2018_paper.html
AUTHORS: Jaeyoung Yoo, Sang-ho Lee, Nojun Kwak
HIGHLIGHT: In this paper, we propose a method to solve the image restoration problem, which tries to restore the details of a corrupted image, especially due to the loss caused by JPEG compression.
- 693, TITLE: Latent RANSAC
http://openaccess.thecvf.com/content_cvpr_2018/html/Korman_Latent_RANSAC_CVPR_2018_paper.html
AUTHORS: Simon Korman, Roe'e Litman
HIGHLIGHT: We present a method that can evaluate a RANSAC hypothesis in constant time, i.e. independent of the size of the data.
- 694, TITLE: Two-Stream Convolutional Networks for Dynamic Texture Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Tesfaldet_Two-Stream_Convolutional_Networks_CVPR_2018_paper.html
AUTHORS: Matthew Tesfaldet, Marcus A. Brubaker, Konstantinos G. Derpanis
HIGHLIGHT: We introduce a two-stream model for dynamic texture synthesis.
- 695, TITLE: Towards Open-Set Identity Preserving Face Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Bao_Towards_Open-Set_Identity_CVPR_2018_paper.html
AUTHORS: Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, Gang Hua
HIGHLIGHT: We propose a framework based on Generative Adversarial Networks to disentangle the identity and attributes of faces, such that we can conveniently recombine different identities and attributes for identity preserving face synthesis in open domains.
- 696, TITLE: A Revised Underwater Image Formation Model
http://openaccess.thecvf.com/content_cvpr_2018/html/Akkaynak_A_Revised_Underwater_CVPR_2018_paper.html
AUTHORS: Derya Akkaynak, Tali Treibitz
HIGHLIGHT: We recently showed that this model introduces significant errors and dependencies in the estimation of the direct transmission signal because underwater, light attenuates in a wavelength-dependent manner.
- 697, TITLE: Graph-Cut RANSAC
http://openaccess.thecvf.com/content_cvpr_2018/html/Barath_Graph-Cut_RANSAC_CVPR_2018_paper.html
AUTHORS: Daniel Barath, Jiří Matas
HIGHLIGHT: A novel method for robust estimation, called Graph-Cut RANSAC, GC-RANSAC in short, is introduced.
- 698, TITLE: Temporal Deformable Residual Networks for Action Segmentation in Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Lei_Temporal_Deformable_Residual_CVPR_2018_paper.html
AUTHORS: Peng Lei, Sinisa Todorovic

HIGHLIGHT: We introduce a new model -- temporal deformable residual network (TDRN) -- aimed at analyzing video intervals at multiple temporal scales for labeling video frames.

699, **TITLE:** Weakly Supervised Action Localization by Sparse Temporal Pooling Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Nguyen_Weakly_Supervised_Action_CVPR_2018_paper.html

AUTHORS: Phuc Nguyen, Ting Liu, Gautam Prasad, Bohyung Han

HIGHLIGHT: We propose a weakly supervised temporal action localization algorithm on untrimmed videos using convolutional neural networks.

700, **TITLE:** PoseFlow: A Deep Motion Representation for Understanding Human Behaviors in Videos

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_PoseFlow_A_Deep_CVPR_2018_paper.html

AUTHORS: Dingwen Zhang, Guangyu Guo, Dong Huang, Junwei Han

HIGHLIGHT: To address this issue, this paper presents a novel deep motion representation, called PoseFlow, which reveals human motion in videos while suppressing background and motion blur, and being robust to occlusion.

701, **TITLE:** FFNet: Video Fast-Forwarding via Reinforcement Learning

http://openaccess.thecvf.com/content_cvpr_2018/html/Lan_FFNet_Video_Fast-Forwarding_CVPR_2018_paper.html

AUTHORS: Shuyue Lan, Rameswar Panda, Qi Zhu, Amit K. Roy-Chowdhury

HIGHLIGHT: In this paper, we introduce FastForwardNet (FFNet), a reinforcement learning agent that gets inspiration from video summarization and does fast-forwarding differently.

702, **TITLE:** Multi-Shot Pedestrian Re-Identification via Sequential Decision Making

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Multi-Shot_Pedestrian_Re-Identification_CVPR_2018_paper.html

AUTHORS: Jianfu Zhang, Naiyan Wang, Liqing Zhang

HIGHLIGHT: In contrary to existing works that aggregate single frames features by time series model such as recurrent neural network, in this paper, we propose an interpretable reinforcement learning based approach to this problem.

703, **TITLE:** Attend and Interact: Higher-Order Object Interactions for Video Understanding

http://openaccess.thecvf.com/content_cvpr_2018/html/Ma_Attend_and_Interact_CVPR_2018_paper.html

AUTHORS: Chih-Yao Ma, Asim Kadav, Iain Melvin, Zsolt Kira, Ghassan AlRegib, Hans Peter Graf

HIGHLIGHT: In this paper, we propose to efficiently learn higher-order interactions between arbitrary subgroups of objects for fine-grained video understanding.

704, **TITLE:** Where and Why Are They Looking? Jointly Inferring Human Attention and Intentions in Complex Tasks

http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Where_and_Why_CVPR_2018_paper.html

AUTHORS: Ping Wei, Yang Liu, Tianmin Shu, Nanning Zheng, Song-Chun Zhu

HIGHLIGHT: We propose a hierarchical model of human-attention-object (HAO) which represents tasks, intentions, and attention under a unified framework.

We built a new video dataset of tasks, intentions, and attention.

705, **TITLE:** Fully Convolutional Adaptation Networks for Semantic Segmentation

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Fully_Convolutional_Adaptation_CVPR_2018_paper.html

AUTHORS: Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, Tao Mei

HIGHLIGHT: In this paper, we facilitate this issue from the perspectives of both visual appearance-level and representation-level domain adaptation.

706, **TITLE:** Semantic Video Segmentation by Gated Recurrent Flow Propagation

http://openaccess.thecvf.com/content_cvpr_2018/html/Nilsson_Semantic_Video_Segmentation_CVPR_2018_paper.html

AUTHORS: David Nilsson, Cristian Sminchisescu

HIGHLIGHT: In this paper we present a deep, end-to-end trainable methodology for video segmentation that is capable of leveraging the information present in unlabeled data, besides sparsely labeled frames, in order to improve semantic estimates.

707, **TITLE:** Interpretable Video Captioning via Trajectory Structured Localization

http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Interpretable_Video_Captioning_CVPR_2018_paper.html

AUTHORS: Xian Wu, Guanbin Li, Qingxing Cao, Qingge Ji, Liang Lin

HIGHLIGHT: In this paper, we propose a Trajectory Structured Attentional Encoder-Decoder (TSA-ED) neural network framework for more elaborate video captioning which works by integrating local spatial-temporal representation at trajectory level through structured attention mechanism.

708, **TITLE:** Deep Hashing via Discrepancy Minimization

http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Deep_Hashing_via_CVPR_2018_paper.html
AUTHORS: Zhixiang Chen, Xin Yuan, Jiwen Lu, Qi Tian, Jie Zhou
HIGHLIGHT: This paper presents a discrepancy minimizing model to address the discrete optimization problem in hashing learning.

709, TITLE: ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_ShuffleNet_An_Extremely_CVPR_2018_paper.html
AUTHORS: Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, Jian Sun
HIGHLIGHT: We introduce an extremely computation-efficient CNN architecture named ShuffleNet, which is designed specially for mobile devices with very limited computing power (e.g., 10-150 MFLOPs).

710, TITLE: Zero-Shot Recognition via Semantic Embeddings and Knowledge Graphs
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Zero-Shot_Recognition_via_CVPR_2018_paper.html
AUTHORS: Xiaolong Wang, Yufei Ye, Abhinav Gupta
HIGHLIGHT: In this paper, we build upon the recently introduced Graph Convolutional Network (GCN) and propose an approach that uses both semantic embeddings and the categorical relationships to predict the classifiers.

711, TITLE: Referring Relationships
http://openaccess.thecvf.com/content_cvpr_2018/html/Krishna_Referring_Relationships_CVPR_2018_paper.html
AUTHORS: Ranjay Krishna, Ines Chami, Michael Bernstein, Li Fei-Fei
HIGHLIGHT: In this paper, we formulate the task of utilizing these "referring relationships" to disambiguate between entities of the same category.

712, TITLE: Improving Object Localization With Fitness NMS and Bounded IoU Loss
http://openaccess.thecvf.com/content_cvpr_2018/html/Tychsen-Smith_Improving_Object_Localization_CVPR_2018_paper.html
AUTHORS: Lachlan Tychsen-Smith, Lars Petersson
HIGHLIGHT: To address this issue we propose a simple and fast modification to the existing methods called Fitness NMS.

713, TITLE: End-to-End Deep Kronecker-Product Matching for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_End-to-End_Deep_Kronecker-Product_CVPR_2018_paper.html
AUTHORS: Yantao Shen, Tong Xiao, Hongsheng Li, Shuai Yi, Xiaogang Wang
HIGHLIGHT: In this paper, we propose a novel Kronecker Product Matching module to match feature maps of different persons in an end-to-end trainable deep neural network.

714, TITLE: Semantic Visual Localization
http://openaccess.thecvf.com/content_cvpr_2018/html/Schonberger_Semantic_Visual_Localization_CVPR_2018_paper.html
AUTHORS: Johannes L. Schönberger, Marc Pollefeys, Andreas Geiger, Torsten Sattler
HIGHLIGHT: In this paper, we propose a novel approach based on a joint 3D geometric and semantic understanding of the world, enabling it to succeed under conditions where previous approaches failed.

715, TITLE: Objects as Context for Detecting Their Semantic Parts
http://openaccess.thecvf.com/content_cvpr_2018/html/Gonzalez-Garcia_Objects_as_Context_CVPR_2018_paper.html
AUTHORS: Abel Gonzalez-García, Davide Modolo, Vittorio Ferrari
HIGHLIGHT: We present a semantic part detection approach that effectively leverages object information.

716, TITLE: End-to-End Weakly-Supervised Semantic Alignment
http://openaccess.thecvf.com/content_cvpr_2018/html/Rocco_End-to-End_Weakly-Supervised_Semantic_CVPR_2018_paper.html
AUTHORS: Ignacio Rocco, Relja Arandjelović, Josef Sivic
HIGHLIGHT: We present the following three principal contributions.

717, TITLE: Dynamic Zoom-In Network for Fast Object Detection in Large Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Gao_Dynamic_Zoom-In_Network_CVPR_2018_paper.html
AUTHORS: Mingfei Gao, Ruichi Yu, Ang Li, Vlad I. Morariu, Larry S. Davis
HIGHLIGHT: We introduce a generic framework that reduces the computational cost of object detection while retaining accuracy for scenarios where objects with varied sizes appear in high resolution images.

718, TITLE: Learning Markov Clustering Networks for Scene Text Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Learning_Markov_Clustering_CVPR_2018_paper.html
AUTHORS: Zichuan Liu, Guosheng Lin, Sheng Yang, Jiashi Feng, Weisi Lin, Wang Ling Goh

- HIGHLIGHT:** A novel framework named Markov Clustering Network (MCN) is proposed for fast and robust scene text detection.
- 719, **TITLE:** Deep Reinforcement Learning of Region Proposal Networks for Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Pirinen_Deep_Reinforcement_Learning_CVPR_2018_paper.html
AUTHORS: Aleksis Pirinen, Cristian Sminchisescu
HIGHLIGHT: We propose drl-RPN, a deep reinforcement learning-based visual recognition model consisting of a sequential region proposal network (RPN) and an object detector.
- 720, **TITLE:** Beyond Holistic Object Recognition: Enriching Image Understanding With Part States
http://openaccess.thecvf.com/content_cvpr_2018/html/Lu_Beyond_Holistic_Object_CVPR_2018_paper.html
AUTHORS: Cewu Lu, Hao Su, Yonglu Li, Yongyi Lu, Li Yi, Chi-Keung Tang, Leonidas J. Guibas
HIGHLIGHT: The other contribution of this paper is our part state dataset which contains rich part-level semantic annotations.
- 721, **TITLE:** Discriminability Objective for Training Descriptive Captions
http://openaccess.thecvf.com/content_cvpr_2018/html/Luo_Discriminability_Objective_for_CVPR_2018_paper.html
AUTHORS: Ruotian Luo, Brian Price, Scott Cohen, Gregory Shakhnarovich
HIGHLIGHT: We propose a way to improve this aspect of caption generation.
- 722, **TITLE:** Visual Question Answering With Memory-Augmented Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Ma_Visual_Question_Answering_CVPR_2018_paper.html
AUTHORS: Chao Ma, Chunhua Shen, Anthony Dick, Qi Wu, Peng Wang, Anton van den Hengel, Ian Reid
HIGHLIGHT: In this paper, we exploit memory-augmented neural networks to predict accurate answers to visual questions, even when those answers rarely occur in the training set.
- 723, **TITLE:** Structure Inference Net: Object Detection Using Scene-Level Context and Instance-Level Relationships
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Structure_Inference_Net_CVPR_2018_paper.html
AUTHORS: Yong Liu, Ruiping Wang, Shiguang Shan, Xilin Chen
HIGHLIGHT: To this end, we present a so-called Structure Inference Network (SIN), a detector that incorporates into a typical detection framework (e.g. Faster R-CNN) with a graphical model which aims to infer object state.
- 724, **TITLE:** Occluded Pedestrian Detection Through Guided Attention in CNNs
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Occluded_Pedestrian_Detection_CVPR_2018_paper.html
AUTHORS: Shanshan Zhang, Jian Yang, Bernt Schiele
HIGHLIGHT: In this paper, we aim to propose a simple and compact method based on the FasterRCNN architecture for occluded pedestrian detection.
- 725, **TITLE:** Reward Learning From Narrated Demonstrations
http://openaccess.thecvf.com/content_cvpr_2018/html/Tung_Reward_Learning_From_CVPR_2018_paper.html
AUTHORS: Hsiao-Yu Tung, Adam W. Harley, Liang-Kang Huang, Katerina Fragkiadaki
HIGHLIGHT: This work proposes joint learning of natural language grounding and instructable behavioural policies reinforced by perceptual detectors of natural language expressions, grounded to the sensory inputs of the robotic agent. We introduce a dataset of NVD where teachers perform activities while describing them in detail.
- 726, **TITLE:** Weakly-Supervised Semantic Segmentation Network With Deep Seeded Region Growing
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Weakly-Supervised_Semantic_Segmentation_CVPR_2018_paper.html
AUTHORS: Zilong Huang, Xinggang Wang, Jiashi Wang, Wenyu Liu, Jingdong Wang
HIGHLIGHT: Inspired by the traditional image segmentation methods of seeded region growing, we propose to train a semantic segmentation network starting from the discriminative regions and progressively increase the pixel-level supervision using by seeded region growing.
- 727, **TITLE:** PoTion: Pose MoTion Representation for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Choutas_PoTion_Pose_MoTion_CVPR_2018_paper.html
AUTHORS: Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, Cordelia Schmid
HIGHLIGHT: In this paper, we claim that considering them jointly offers rich information for action recognition.
- 728, **TITLE:** Bilateral Ordinal Relevance Multi-Instance Regression for Facial Action Unit Intensity Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Bilateral_Ordinal_Relevance_CVPR_2018_paper.html
AUTHORS: Yong Zhang, Rui Zhao, Weiming Dong, Bao-Gang Hu, Qiang Ji

HIGHLIGHT: In this paper, we propose a novel weakly supervised regression model-Bilateral Ordinal Relevance Multi-instance Regression (BORMIR), which learns a frame-level intensity estimator with weakly labeled sequences.

729, **TITLE:** Pulling Actions out of Context: Explicit Separation for Effective Combination
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Pulling_Actions_out_CVPR_2018_paper.html

AUTHORS: Yang Wang, Minh Hoai

HIGHLIGHT: In this paper, we propose a novel approach for training a human action recognizer, one that can: (1) explicitly factorize human actions from the co-occurring factors; (2) deliberately build a model for human actions and a separate model for all correlated contextual elements; and (3) effectively combine the models for human action recognition.

730, **TITLE:** Dynamic Feature Learning for Partial Face Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/He_Dynamic_Feature_Learning_CVPR_2018_paper.html

AUTHORS: Lingxiao He, Haiqing Li, Qi Zhang, Zhenan Sun

HIGHLIGHT: This study combines Fully Convolutional Network (FCN) with Sparse Representation Classification (SRC) to propose a novel partial face recognition approach, called Dynamic Feature Matching (DFM), to address partial face images regardless of sizes.

731, **TITLE:** Exploiting Transitivity for Learning Person Re-Identification Models on a Budget
http://openaccess.thecvf.com/content_cvpr_2018/html/Roy_Exploiting_Transitivity_for_CVPR_2018_paper.html

AUTHORS: Sourya Roy, Sujoy Paul, Neal E. Young, Amit K. Roy-Chowdhury

HIGHLIGHT: In this work, we focus on this labeling effort minimization problem and approach it as a subset selection task where the objective is to select an optimal subset of image-pairs for labeling without compromising performance.

732, **TITLE:** Deep Spatial Feature Reconstruction for Partial Person Re-Identification: Alignment-Free Approach
http://openaccess.thecvf.com/content_cvpr_2018/html/He_Deep_Spatial_Feature_Reconstruction_CVPR_2018_paper.html

AUTHORS: Lingxiao He, Jian Liang, Haiqing Li, Zhenan Sun

HIGHLIGHT: In this paper, we propose a fast and accurate matching method to address this problem.

733, **TITLE:** Every Smile Is Unique: Landmark-Guided Diverse Smile Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Every_Smile_Is_CVPR_2018_paper.html

AUTHORS: Wei Wang, Xavier Alameda-Pineda, Dan Xu, Pascal Fua, Elisa Ricci, Nicu Sebe

HIGHLIGHT: To tackle this one-to-many video generation problem, we propose a novel deep learning architecture named Conditional Multi-Mode Network (CMM-Net).

734, **TITLE:** UV-GAN: Adversarial Facial UV Map Completion for Pose-Invariant Face Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Deng_UV-GAN_Adversarial_Facial_CVPR_2018_paper.html

AUTHORS: Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, Stefanos Zafeiriou

HIGHLIGHT: In this paper, we propose a framework for training Deep Convolutional Neural Network (DCNN) to complete the facial UV map extracted from in-the-wild images.

We will release the first in-the-wild UV dataset (we refer as WildUV) that comprises of complete facial UV maps from 1,892 identities for research purposes.

735, **TITLE:** Cascaded Pyramid Network for Multi-Person Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Cascaded_Pyramid_Network_CVPR_2018_paper.html

AUTHORS: Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, Jian Sun

HIGHLIGHT: In this paper, we present a novel network structure called Cascaded Pyramid Network (CPN) which targets to relieve the problem from these "hard" keypoints. More specifically, our algorithm includes two stages: GlobalNet and RefineNet.

736, **TITLE:** A Face-to-Face Neural Conversation Model
http://openaccess.thecvf.com/content_cvpr_2018/html/Chu_A_Face-to-Face_Neural_CVPR_2018_paper.html

AUTHORS: Hang Chu, Daiqing Li, Sanja Fidler

HIGHLIGHT: We propose a neural conversation model that aims to read and generate facial gestures alongside with text.

737, **TITLE:** End-to-End Recovery of Human Shape and Pose
http://openaccess.thecvf.com/content_cvpr_2018/html/Kanazawa_End-to-End_Recovery_of_CVPR_2018_paper.html

AUTHORS: Angjoo Kanazawa, Michael J. Black, David W. Jacobs, Jitendra Malik

HIGHLIGHT: We describe Human Mesh Recovery (HMR), an end-to-end framework for reconstructing a full 3D mesh of a human body from a single RGB image.

738, **TITLE:** Squeeze-and-Excitation Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Squeeze-and-Excitation_Networks_CVPR_2018_paper.html
AUTHORS: Jie Hu, Li Shen, Gang Sun
HIGHLIGHT: In this work, we focus on the channel relationship and propose a novel architectural unit, which we term the “Squeeze-and-Excitation” (SE) block, that adaptively recalibrates channel-wise feature responses by explicitly modelling interdependencies between channels.

739, TITLE: Revisiting Salient Object Detection: Simultaneous Detection, Ranking, and Subitizing of Multiple Salient Objects
http://openaccess.thecvf.com/content_cvpr_2018/html/Islam_Revisiting_Salient_Object_CVPR_2018_paper.html
AUTHORS: Md Amirul Islam, Mahmoud Kalash, Neil D. B. Bruce
HIGHLIGHT: In this paper, we argue that work to date has addressed a problem that is relatively ill-posed. The solution presented in this paper solves this more general problem that considers relative rank, and we propose data and metrics suitable to measuring success in a relative object saliency landscape.

740, TITLE: Context Encoding for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Context_Encoding_for_CVPR_2018_paper.html
AUTHORS: Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Amrith Tyagi, Amit Agrawal
HIGHLIGHT: In this paper, we explore the impact of global contextual information in semantic segmentation by introducing the Context Encoding Module, which captures the semantic context of scenes and selectively highlights class-dependent featuremaps.

741, TITLE: Creating Capsule Wardrobes From Fashion Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Hsiao_Creating_Capsule_Wardrobes_CVPR_2018_paper.html
AUTHORS: Wei-Lin Hsiao, Kristen Grauman
HIGHLIGHT: We propose to automatically create $\text{emph}\{\text{capsule wardrobes}\}$.

742, TITLE: Webly Supervised Learning Meets Zero-Shot Learning: A Hybrid Approach for Fine-Grained Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Niu_Webly_Supervised_Learning_CVPR_2018_paper.html
AUTHORS: Li Niu, Ashok Veeraraghavan, Ashutosh Sabharwal
HIGHLIGHT: The drawbacks of the above two directions motivate us to design a new framework which can jointly leverage both web data and auxiliary labeled categories to predict the test categories that are not associated with any well-labeled training images.

743, TITLE: Look, Imagine and Match: Improving Textual-Visual Cross-Modal Retrieval With Generative Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Gu_Look_Imagine_and_CVPR_2018_paper.html
AUTHORS: Jiuxiang Gu, Jianfei Cai, Shafiq R. Joty, Li Niu, Gang Wang
HIGHLIGHT: Unlike existing image-text retrieval approaches that embed image-text pairs as single feature vectors in a common representational space, we propose to incorporate generative processes into the cross-modal feature embedding, through which we are able to learn not only the global abstract features but also the local grounded features.

744, TITLE: Bidirectional Attentive Fusion With Context Gating for Dense Video Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Bidirectional_Attentive_Fusion_CVPR_2018_paper.html
AUTHORS: Jingwen Wang, Wenhao Jiang, Lin Ma, Wei Liu, Yong Xu
HIGHLIGHT: We propose a bidirectional proposal method that effectively exploits both past and future contexts to make proposal predictions.

745, TITLE: InLoc: Indoor Visual Localization With Dense Matching and View Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Taira_InLoc_Indoor_Visual_CVPR_2018_paper.html
AUTHORS: Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, Akihiko Torii
HIGHLIGHT: The contributions of this work are three-fold. Second, we collect a new dataset with reference 6DoF poses for large-scale indoor localization.

746, TITLE: Towards High Performance Video Object Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_Towards_High_Performance_CVPR_2018_paper.html
AUTHORS: Xizhou Zhu, Jifeng Dai, Lu Yuan, Yichen Wei
HIGHLIGHT: Built upon the recent works, this work proposes a unified viewpoint based on the principle of multi-frame end-to-end learning of features and cross-frame motion.

747, TITLE: Neural Baby Talk
http://openaccess.thecvf.com/content_cvpr_2018/html/Lu_Neural_Baby_Talk_CVPR_2018_paper.html
AUTHORS: Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh

HIGHLIGHT: We introduce a novel framework for image captioning that can produce natural language explicitly grounded in entities that object detectors find in the image.

748, **TITLE:** Few-Shot Image Recognition by Predicting Parameters From Activations
http://openaccess.thecvf.com/content_cvpr_2018/html/Qiao_Few-Shot_Image_Recognition_CVPR_2018_paper.html
AUTHORS: Siyuan Qiao, Chenxi Liu, Wei Shen, Alan L. Yuille
HIGHLIGHT: In this paper, we are interested in the few-shot learning problem.

749, **TITLE:** Iterative Visual Reasoning Beyond Convolutions
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Iterative_Visual_Reasoning_CVPR_2018_paper.html
AUTHORS: Xinlei Chen, Li-Jia Li, Li Fei-Fei, Abhinav Gupta
HIGHLIGHT: We present a novel framework for iterative visual reasoning.

750, **TITLE:** Visual Question Reasoning on General Dependency Tree
http://openaccess.thecvf.com/content_cvpr_2018/html/Cao_Visual_Question_Reasoning_CVPR_2018_paper.html
AUTHORS: Qingxing Cao, Xiaodan Liang, Bailing Li, Guanbin Li, Liang Lin
HIGHLIGHT: In this paper, to enable global context reasoning for better aligning image and language domains in diverse and unrestricted cases, we propose a novel reasoning network called Adversarial Composition Modular Network (ACMN).

751, **TITLE:** CVM-Net: Cross-View Matching Network for Image-Based Ground-to-Aerial Geo-Localization
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_CVM-Net_Cross-View_Matching_CVPR_2018_paper.html
AUTHORS: Sixing Hu, Mengdan Feng, Rang M. H. Nguyen, Gim Hee Lee
HIGHLIGHT: We leverage on the recent success of deep learning to propose the CVM-Net for the cross-view image-based ground-to-aerial geo-localization task.

752, **TITLE:** Revisiting Dilated Convolution: A Simple Approach for Weakly- and Semi-Supervised Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Revisiting_Dilated_Convolution_CVPR_2018_paper.html
AUTHORS: Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, Thomas S. Huang
HIGHLIGHT: In this work, we revisit the dilated convolution proposed in [1] and shed light on how it enables the classification network to generate dense object localization.

753, **TITLE:** Low-Shot Learning From Imaginary Data
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Low-Shot_Learning_From_CVPR_2018_paper.html
AUTHORS: Yu-Xiong Wang, Ross Girshick, Martial Hebert, Bharath Hariharan
HIGHLIGHT: We present a novel approach to low-shot learning that uses this idea.

754, **TITLE:** DoubleFusion: Real-Time Capture of Human Performances With Inner Body Shapes From a Single Depth Sensor
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_DoubleFusion_Real-Time_Capture_CVPR_2018_paper.html
AUTHORS: Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, Yebin Liu
HIGHLIGHT: We propose DoubleFusion, a new real-time system that combines volumetric dynamic reconstruction with data-driven template fitting to simultaneously reconstruct detailed geometry, non-rigid motion and the inner human body shape from a single depth camera.

755, **TITLE:** DensePose: Dense Human Pose Estimation in the Wild
http://openaccess.thecvf.com/content_cvpr_2018/html/Guler_DensePose_Dense_Human_CVPR_2018_paper.html
AUTHORS: R?za Alp Güler, Natalia Neverova, Iasonas Kokkinos
HIGHLIGHT: In this work we establish dense correspondences between an RGB image and a surface-based representation of the human body, a task we refer to as dense human pose estimation.

756, **TITLE:** Ordinal Depth Supervision for 3D Human Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Pavlakos_Ordinal_Depth_Supervision_CVPR_2018_paper.html
AUTHORS: Georgios Pavlakos, Xiaowei Zhou, Kostas Daniilidis
HIGHLIGHT: To alleviate the need for accurate 3D ground truth, we propose to use a weaker supervision signal provided by the ordinal depths of human joints.

757, **TITLE:** Consensus Maximization for Semantic Region Correspondences
http://openaccess.thecvf.com/content_cvpr_2018/html/Speciale_Consensus_Maximization_for_CVPR_2018_paper.html
AUTHORS: Pablo Speciale, Danda P. Paudel, Martin R. Oswald, Hayko Riemenschneider, Luc Van Gool, Marc Pollefeys
HIGHLIGHT: We propose a novel method for the geometric registration of semantically labeled regions.

- 758, TITLE: Robust Hough Transform Based 3D Reconstruction From Circular Light Fields
http://openaccess.thecvf.com/content_cvpr_2018/html/Vianello_Robust_Hough_Transform_CVPR_2018_paper.html
AUTHORS: Alessandro Vianello, Jens Ackermann, Maximilian Diebold, Bernd Jähne
HIGHLIGHT: This paper presents a novel method which allows to reconstruct depth information from data acquired with a circular camera motion, termed circular light fields.
- 759, TITLE: Alive Caricature From 2D to 3D
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Alive_Caricature_From_CVPR_2018_paper.html
AUTHORS: Qianyi Wu, Juyong Zhang, Yu-Kun Lai, Jianmin Zheng, Jianfei Cai
HIGHLIGHT: Built upon the proposed deformation representation, an optimization model is formulated to find the 3D caricature that captures the style of the 2D caricature image automatically.
- 760, TITLE: Nonlinear 3D Face Morphable Model
http://openaccess.thecvf.com/content_cvpr_2018/html/Tran_Nonlinear_3D_Face_CVPR_2018_paper.html
AUTHORS: Luan Tran, Xiaoming Liu
HIGHLIGHT: To address these problems, this paper proposes an innovative framework to learn a nonlinear 3DMM model from a large set of unconstrained face images, without collecting 3D face scans.
- 761, TITLE: Through-Wall Human Pose Estimation Using Radio Signals
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_Through-Wall_Human_Pose_CVPR_2018_paper.html
AUTHORS: Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, Dina Katabi
HIGHLIGHT: We introduce a deep neural network approach that parses such radio signals to estimate 2D poses.
- 762, TITLE: What Makes a Video a Video: Analyzing Temporal Information in Video Understanding Models and Datasets
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_What_Makes_a_CVPR_2018_paper.html
AUTHORS: De-An Huang, Vignesh Ramanathan, Dhruv Mahajan, Lorenzo Torresani, Manohar Paluri, Li Fei-Fei, Juan Carlos Niebles
HIGHLIGHT: In this work, we aim to bridge this gap and ask the following question: How important is the motion in the video for recognizing the action?
- 763, TITLE: Fast Video Object Segmentation by Reference-Guided Mask Propagation
http://openaccess.thecvf.com/content_cvpr_2018/html/Oh_Fast_Video_Object_CVPR_2018_paper.html
AUTHORS: Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, Seon Joo Kim
HIGHLIGHT: We present an efficient method for the semi-supervised video object segmentation.
- 764, TITLE: NeuralNetwork-Viterbi: A Framework for Weakly Supervised Video Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Richard_NeuralNetwork-Viterbi_A_Framework_CVPR_2018_paper.html
AUTHORS: Alexander Richard, Hilde Kuehne, Ahsan Iqbal, Juergen Gall
HIGHLIGHT: In this work, we propose a novel learning algorithm with a Viterbi-based loss that allows for online and incremental learning of weakly annotated video data.
- 765, TITLE: Actor and Observer: Joint Modeling of First and Third-Person Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Sigurdsson_Actor_and_Observer_CVPR_2018_paper.html
AUTHORS: Gunnar A. Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, Karteek Alahari
HIGHLIGHT: Several theories in cognitive neuroscience suggest that when people interact with the world, or simulate interactions, they do so from a first-person egocentric perspective, and seamlessly transfer knowledge between third-person (observer) and first-person (actor).
- 766, TITLE: HSA-RNN: Hierarchical Structure-Adaptive RNN for Video Summarization
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhao_HSA-RNN_Hierarchical_Structure-Adaptive_CVPR_2018_paper.html
AUTHORS: Bin Zhao, Xuelong Li, Xiaoqiang Lu
HIGHLIGHT: To address this problem, we propose a structure-adaptive video summarization approach that integrates shot segmentation and video summarization into a Hierarchical Structure-Adaptive RNN, denoted as HSA-RNN.
- 767, TITLE: Fast and Accurate Online Video Object Segmentation via Tracking Parts
http://openaccess.thecvf.com/content_cvpr_2018/html/Cheng_Fast_and_Accurate_CVPR_2018_paper.html
AUTHORS: Jingchun Cheng, Yi-Hsuan Tsai, Wei-Chih Hung, Shengjin Wang, Ming-Hsuan Yang

HIGHLIGHT: In this paper, we propose a fast and accurate video object segmentation algorithm that can immediately start the segmentation process once receiving the images.

768, **TITLE:** Now You Shake Me: Towards Automatic 4D Cinema
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Now_You_Shake_CVPR_2018_paper.html

AUTHORS: Yuhao Zhou, Makarand Tapaswi, Sanja Fidler

HIGHLIGHT: We propose a Conditional Random Field model atop a neural network that brings together visual and audio information, as well as semantics in the form of person tracks.
We collect a new dataset referred to as the Movie4D dataset which annotates over 9K effects in 63 movies.

769, **TITLE:** Viewpoint-Aware Video Summarization
http://openaccess.thecvf.com/content_cvpr_2018/html/Kanehira_Viewpoint-Aware_Video_Summarization_CVPR_2018_paper.html

AUTHORS: Atsushi Kanehira, Luc Van Gool, Yoshitaka Ushiku, Tatsuya Harada

HIGHLIGHT: To satisfy these requirements (A)-(C) simultaneously, we proposed a novel video summarization method from multiple groups of videos.

Moreover, we developed a novel dataset to investigate how well the generated summary reflects the underlying viewpoint.

770, **TITLE:** Photometric Stereo in Participating Media Considering Shape-Dependent Forward Scatter
http://openaccess.thecvf.com/content_cvpr_2018/html/Fujimura_Photometric_Stereo_in_CVPR_2018_paper.html

AUTHORS: Yuki Fujimura, Masaaki Iiyama, Atsushi Hashimoto, Michihiko Minoh

HIGHLIGHT: In this paper, we propose a photometric stereo method for participating media.

771, **TITLE:** Direction-Aware Spatial Context Features for Shadow Detection
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Direction-Aware_Spatial_Context_CVPR_2018_paper.html

AUTHORS: Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, Pheng-Ann Heng

HIGHLIGHT: This paper presents a novel network for shadow detection by analyzing image context in a direction-aware manner.

772, **TITLE:** Discriminative Learning of Latent Features for Zero-Shot Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Discriminative_Learning_of_CVPR_2018_paper.html

AUTHORS: Yan Li, Junge Zhang, Jianguo Zhang, Kaiqi Huang

HIGHLIGHT: In this work, we retrospect existing methods and demonstrate the necessity to learn discriminative representations for both visual and semantic instances of ZSL.

773, **TITLE:** Learning to Adapt Structured Output Space for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Tsai_Learning_to_Adapt_CVPR_2018_paper.html

AUTHORS: Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, Manmohan Chandraker

HIGHLIGHT: In this paper, we propose an adversarial learning method for domain adaptation in the context of semantic segmentation.

774, **TITLE:** Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics
http://openaccess.thecvf.com/content_cvpr_2018/html/Kendall_Multi-Task_Learning_Using_CVPR_2018_paper.html

AUTHORS: Alex Kendall, Yarin Gal, Roberto Cipolla

HIGHLIGHT: In this paper we make the observation that the performance of such systems is strongly dependent on the relative weighting between each task's loss.

775, **TITLE:** Jointly Localizing and Describing Events for Dense Video Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Jointly_Localizing_and_CVPR_2018_paper.html

AUTHORS: Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, Tao Mei

HIGHLIGHT: In this paper, we present a novel framework for dense video captioning that unifies the localization of temporal event proposals and sentence generation of each proposal, by jointly training them in an end-to-end manner.

776, **TITLE:** Going From Image to Video Saliency: Augmenting Image Saliency With Dynamic Attentional Push
http://openaccess.thecvf.com/content_cvpr_2018/html/Gorji_Going_From_Image_CVPR_2018_paper.html

AUTHORS: Siavash Gorji, James J. Clark

HIGHLIGHT: We present a novel method to incorporate the recent advent in static saliency models to predict the saliency in videos.

777, **TITLE:** M3: Multimodal Memory Modelling for Video Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_M3_Multimodal_Memory_CVPR_2018_paper.html

AUTHORS: Junbo Wang, Wei Wang, Yan Huang, Liang Wang, Tieniu Tan
HIGHLIGHT: Inspired by the facts that memory modelling poses potential advantages to long-term sequential problems [35] and working memory is the key factor of visual attention [33], we propose a Multimodal Memory Model (M3) to describe videos, which builds a visual and textual shared memory to model the long-term visual-textual dependency and further guide visual attention on described visual targets to solve visual-textual alignments.

778, **TITLE:** Emotional Attention: A Study of Image Sentiment and Visual Attention
http://openaccess.thecvf.com/content_cvpr_2018/html/Fan_Emotional_Attention_A_CVPR_2018_paper.html
AUTHORS: Shaojing Fan, Zhiqi Shen, Ming Jiang, Bryan L. Koenig, Juan Xu, Mohan S. Kankanhalli, Qi Zhao
HIGHLIGHT: In this paper, we present the first study to focus on the relation between emotional properties of an image and visual attention.
We first create the EMOTional attention dataset (EMOD).

779, **TITLE:** A Low Power, High Throughput, Fully Event-Based Stereo System
http://openaccess.thecvf.com/content_cvpr_2018/html/Andreopoulos_A_Low_Power_CVPR_2018_paper.html
AUTHORS: Alexander Andreopoulos, Hiram J. Kashyap, Tapan K. Nayak, Arnon Amir, Myron D. Flickner
HIGHLIGHT: We introduce a stereo correspondence system implemented fully on event-based digital hardware, using a fully graph-based non von-Neumann computation model, where no frames, arrays, or any other such data-structures are used.

780, **TITLE:** VITON: An Image-Based Virtual Try-On Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Han_VITON_An_Image-Based_CVPR_2018_paper.html
AUTHORS: Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, Larry S. Davis
HIGHLIGHT: We present an image-based Virtual Try-On Network (VITON) without using 3D information in any form, which seamlessly transfers a desired clothing item onto the corresponding region of a person using a coarse-to-fine strategy.

781, **TITLE:** Multi-Oriented Scene Text Detection via Corner Localization and Region Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Lyu_Multi-Oriented_Scene_Text_CVPR_2018_paper.html
AUTHORS: Pengyuan Lyu, Cong Yao, Wenhao Wu, Shuicheng Yan, Xiang Bai
HIGHLIGHT: In this paper, we present a method that combines the ideas of the two types of methods while avoiding their shortcomings.

782, **TITLE:** Multi-Content GAN for Few-Shot Font Style Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Azadi_Multi-Content_GAN_for_CVPR_2018_paper.html
AUTHORS: Samaneh Azadi, Matthew Fisher, Vladimir G. Kim, Zhaowen Wang, Eli Shechtman, Trevor Darrell
HIGHLIGHT: In this work, we focus on the challenge of taking partial observations of highly-stylized text and generalizing the observations to generate unobserved glyphs in the ornamented typeface.

783, **TITLE:** Audio to Body Dynamics
http://openaccess.thecvf.com/content_cvpr_2018/html/Shlizerman_Audio_to_Body_CVPR_2018_paper.html
AUTHORS: Eli Shlizerman, Lucio Dery, Hayden Schoen, Ira Kemelmacher-Shlizerman
HIGHLIGHT: We present a method that gets as input an audio of violin or piano playing, and outputs a video of skeleton predictions which are further used to animate an avatar.

784, **TITLE:** Weakly Supervised Coupled Networks for Visual Sentiment Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_Weakly_Supervised_Coupled_CVPR_2018_paper.html
AUTHORS: Jufeng Yang, Dongyu She, Yu-Kun Lai, Paul L. Rosin, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we solve the problem of visual sentiment analysis using the high-level abstraction in the recognition process.

785, **TITLE:** Future Person Localization in First-Person Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Yagi_Future_Person_Localization_CVPR_2018_paper.html
AUTHORS: Takuma Yagi, Kartikeya Mangalam, Ryo Yonetani, Yoichi Sato
HIGHLIGHT: We present a new task that predicts future locations of people observed in first-person videos.

786, **TITLE:** Preserving Semantic Relations for Zero-Shot Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Annadani_Preserving_Semantic_Relations_CVPR_2018_paper.html
AUTHORS: Yashas Annadani, Soma Biswas
HIGHLIGHT: In this work, we propose to utilize the structure of the space spanned by the attributes using a set of relations.

787, **TITLE:** Show Me a Story: Towards Coherent Neural Story Illustration

- http://openaccess.thecvf.com/content_cvpr_2018/html/Ravi_Show_Me_a_CVPR_2018_paper.html
AUTHORS: Hareesh Ravi, Lezi Wang, Carlos Muniz, Leonid Sigal, Dimitris Metaxas, Mubbasir Kapadia
HIGHLIGHT: We propose an end-to-end network for the visual illustration of a sequence of sentences forming a story.
- 788, TITLE: Reconstruction Network for Video Captioning
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Reconstruction_Network_for_CVPR_2018_paper.html
AUTHORS: Bairui Wang, Lin Ma, Wei Zhang, Wei Liu
HIGHLIGHT: In this paper, the problem of describing visual contents of a video sequence with natural language is addressed.
- 789, TITLE: Fast Spectral Ranking for Similarity Search
http://openaccess.thecvf.com/content_cvpr_2018/html/Iscen_Fast_Spectral_Ranking_CVPR_2018_paper.html
AUTHORS: Ahmet Iscen, Yannis Avrithis, Giorgos Toliias, Teddy Furon, Ondrej Chum
HIGHLIGHT: This work introduces an explicit embedding reducing manifold search to Euclidean search followed by dot product similarity search.
- 790, TITLE: Mining on Manifolds: Metric Learning Without Labels
http://openaccess.thecvf.com/content_cvpr_2018/html/Iscen_Mining_on_Manifolds_CVPR_2018_paper.html
AUTHORS: Ahmet Iscen, Giorgos Toliias, Yannis Avrithis, Ondrej Chum
HIGHLIGHT: In this work we present a novel unsupervised framework for hard training example mining.
- 791, TITLE: PIXOR: Real-Time 3D Object Detection From Point Clouds
http://openaccess.thecvf.com/content_cvpr_2018/html/Yang_PIXOR_Real-Time_3D_CVPR_2018_paper.html
AUTHORS: Bin Yang, Wenjie Luo, Raquel Urtasun
HIGHLIGHT: We utilize the 3D data more efficiently by representing the scene from the Bird's Eye View (BEV), and propose PIXOR, a proposal-free, single-stage detector that outputs oriented 3D object estimates decoded from pixel-wise neural network predictions.
- 792, TITLE: Leveraging Unlabeled Data for Crowd Counting by Learning to Rank
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Leveraging_Unlabeled_Data_CVPR_2018_paper.html
AUTHORS: Xialei Liu, Joost van de Weijer, Andrew D. Bagdanov
HIGHLIGHT: We propose a novel crowd counting approach that leverages abundantly available unlabeled crowd imagery in a learning-to-rank framework.
We collect two crowd scene datasets from Google using keyword searches and query-by-example image retrieval, respectively.
- 793, TITLE: Zero-Shot Kernel Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Zero-Shot_Kernel_Learning_CVPR_2018_paper.html
AUTHORS: Hongguang Zhang, Piotr Koniusz
HIGHLIGHT: In this paper, we address an open problem of zero-shot learning.
- 794, TITLE: Differential Attention for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Patro_Differential_Attention_for_CVPR_2018_paper.html
AUTHORS: Badri Patro, Vinay P. Namboodiri
HIGHLIGHT: In this paper we aim to answer questions based on images when provided with a dataset of question-answer pairs for a number of images during training.
- 795, TITLE: Learning From Noisy Web Data With Category-Level Supervision
http://openaccess.thecvf.com/content_cvpr_2018/html/Niu_Learning_From_Noisy_CVPR_2018_paper.html
AUTHORS: Li Niu, Qingtao Tang, Ashok Veeraraghavan, Ashutosh Sabharwal
HIGHLIGHT: Instead, we propose to address the label noise by using more accessible category-level supervision.
- 796, TITLE: Toward Driving Scene Understanding: A Dataset for Learning Driver Behavior and Causal Reasoning
http://openaccess.thecvf.com/content_cvpr_2018/html/Ramanishka_Toward_Driving_Scene_CVPR_2018_paper.html
AUTHORS: Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, Kate Saenko
HIGHLIGHT: We present the Honda Research Institute Driving Dataset (HDD), a challenging dataset to enable research on learning driver behavior in real-life environments.
- 797, TITLE: Learning Attribute Representations With Localization for Flexible Fashion Search
http://openaccess.thecvf.com/content_cvpr_2018/html/Ak_Learning_Attribute_Representations_CVPR_2018_paper.html
AUTHORS: Kenan E. Ak, Ashraf A. Kassim, Joo Hwee Lim, Jo Yew Tham
HIGHLIGHT: In this paper, we investigate ways of conducting a detailed fashion search using query images and attributes.

798, TITLE: Bidirectional Retrieval Made Simple
http://openaccess.thecvf.com/content_cvpr_2018/html/Wehrmann_Bidirectional_Retrieval_Made_CVPR_2018_paper.html
AUTHORS: Jónatas Wehrmann, Rodrigo C. Barros
HIGHLIGHT: We introduce an efficient character-level inception module, designed to learn textual semantic embeddings by convolving raw characters in distinct granularity levels.

799, TITLE: Learning Multi-Instance Enriched Image Representations via Non-Greedy Ratio Maximization of the l1-Norm Distances
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Learning_Multi-Instance_Enriched_CVPR_2018_paper.html
AUTHORS: Kai Liu, Hua Wang, Feiping Nie, Hao Zhang
HIGHLIGHT: To tackle these two challenges, in this paper we propose a novel image representation learning method that can integrate the local patches (the instances) of an input image (the bag) and its holistic representation into one single-vector representation.

800, TITLE: Learning Visual Knowledge Memory Networks for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Su_Learning_Visual_Knowledge_CVPR_2018_paper.html
AUTHORS: Zhou Su, Chen Zhu, Yinpeng Dong, Dongqi Cai, Yurong Chen, Jianguo Li
HIGHLIGHT: This paper proposes visual knowledge memory network (VKMN) to address this issue, which seamlessly incorporates structured human knowledge and deep visual features into memory networks in an end-to-end learning framework.

801, TITLE: Visual Grounding via Accumulated Attention
http://openaccess.thecvf.com/content_cvpr_2018/html/Deng_Visual_Grounding_via_CVPR_2018_paper.html
AUTHORS: Chaorui Deng, Qi Wu, Qingyao Wu, Fuyuan Hu, Fan Lyu, Mingkui Tan
HIGHLIGHT: In this paper, we formulate these challenges as three attention problems and propose an accumulated attention (A-ATT) mechanism to reason among them jointly.

802, TITLE: Beyond Trade-Off: Accelerate FCN-Based Face Detector With Higher Accuracy
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_Beyond_Trade-Off_Accelerate_CVPR_2018_paper.html
AUTHORS: Guanglu Song, Yu Liu, Ming Jiang, Yujie Wang, Junjie Yan, Biao Leng
HIGHLIGHT: Based on this philosophy, a novel method named scale estimation and spatial attention proposal (S²AP) is proposed to pay attention to some specific scales in image pyramid and valid locations in each scales layer.

803, TITLE: PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning
http://openaccess.thecvf.com/content_cvpr_2018/html/Mallya_PackNet_Adding_Multiple_CVPR_2018_paper.html
AUTHORS: Arun Mallya, Svetlana Lazebnik
HIGHLIGHT: This paper presents a method for adding multiple tasks to a single deep neural network while avoiding catastrophic forgetting.

804, TITLE: Repulsion Loss: Detecting Pedestrians in a Crowd
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Repulsion_Loss_Detecting_CVPR_2018_paper.html
AUTHORS: Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, Chunhua Shen
HIGHLIGHT: Repulsion Loss: Detecting Pedestrians in a Crowd

805, TITLE: Neural Sign Language Translation
http://openaccess.thecvf.com/content_cvpr_2018/html/Camgoz_Neural_Sign_Language_CVPR_2018_paper.html
AUTHORS: Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, Richard Bowden
HIGHLIGHT: In contrast, we introduce the Sign Language Translation (SLT) problem.
To evaluate the performance of Neural SLT, we collected the first publicly available Continuous SLT dataset, RWTH-PHOENIX-Weather 2014T.

806, TITLE: Non-Local Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Non-Local_Neural_Networks_CVPR_2018_paper.html
AUTHORS: Xiaolong Wang, Ross Girshick, Abhinav Gupta, Kaiming He
HIGHLIGHT: In this paper, we present non-local operations as a generic family of building blocks for capturing long-range dependencies.

807, TITLE: LAMV: Learning to Align and Match Videos With Kernelized Temporal Layers
http://openaccess.thecvf.com/content_cvpr_2018/html/Baraldi_LAMV_Learning_to_CVPR_2018_paper.html
AUTHORS: Lorenzo Baraldi, Matthijs Douze, Rita Cucchiara, Hervé Jégou

HIGHLIGHT: Our architecture builds upon and revisits temporal match kernels within neural networks: we propose a new temporal layer that finds temporal alignments by maximizing the scores between two sequences of vectors, according to a time-sensitive similarity metric parametrized in the Fourier domain.

808, TITLE: Optimizing Video Object Detection via a Scale-Time Lattice
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Optimizing_Video_Object_CVPR_2018_paper.html
AUTHORS: Kai Chen, Jiaqi Wang, Shuo Yang, Xingcheng Zhang, Yuanjun Xiong, Chen Change Loy, Dahua Lin
HIGHLIGHT: Specifically, we present a unified framework that integrates detection, temporal propagation, and across-scale refinement on a Scale-Time Lattice.

809, TITLE: Learning Compressible 360° Video Isomers
http://openaccess.thecvf.com/content_cvpr_2018/html/Su_Learning_Compressible_360deg_CVPR_2018_paper.html
AUTHORS: Yu-Chuan Su, Kristen Grauman
HIGHLIGHT: We introduce an approach to predict the sphere rotation that will yield the maximal compression rate.

810, TITLE: Attention Clusters: Purely Attention Based Local Feature Integration for Video Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Long_Attention_Clusters_Purely_CVPR_2018_paper.html
AUTHORS: Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, Shilei Wen
HIGHLIGHT: Accounting for the characteristics of such features in video classification, we propose a local feature integration framework based on attention clusters, and introduce a shifting operation to capture more diverse signals.

811, TITLE: What Have We Learned From Deep Representations for Action Recognition?
http://openaccess.thecvf.com/content_cvpr_2018/html/Feichtenhofer_What_Have_We_CVPR_2018_paper.html
AUTHORS: Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes, Andrew Zisserman
HIGHLIGHT: In this paper, we shed light on deep spatiotemporal representations by visualizing what two-stream models have learned in order to recognize actions in video.

812, TITLE: Controllable Video Generation With Sparse Trajectories
http://openaccess.thecvf.com/content_cvpr_2018/html/Hao_Controllable_Video_Generation_CVPR_2018_paper.html
AUTHORS: Zekun Hao, Xun Huang, Serge Belongie
HIGHLIGHT: In this work, we present a conditional video generation model that allows detailed control over the motion of the generated video.

813, TITLE: Representing and Learning High Dimensional Data With the Optimal Transport Map From a Probabilistic Viewpoint
http://openaccess.thecvf.com/content_cvpr_2018/html/Park_Representing_and_Learning_CVPR_2018_paper.html
AUTHORS: Serim Park, Matthew Thorpe
HIGHLIGHT: In this paper, we propose a generative model in the space of diffeomorphic deformation maps.

814, TITLE: CLIP-Q: Deep Network Compression Learning by In-Parallel Pruning-Quantization
http://openaccess.thecvf.com/content_cvpr_2018/html/Tung_CLIP-Q_Deep_Network_CVPR_2018_paper.html
AUTHORS: Frederick Tung, Greg Mori
HIGHLIGHT: In this paper, we combine network pruning and weight quantization in a single learning framework that performs pruning and quantization jointly, and in parallel with fine-tuning.

815, TITLE: Inference in Higher Order MRF-MAP Problems With Small and Large Cliques
http://openaccess.thecvf.com/content_cvpr_2018/html/Shanu_Inference_in_Higher_CVPR_2018_paper.html
AUTHORS: Ishant Shanu, Chetan Arora, S.N. Maheshwari
HIGHLIGHT: We show in this paper that the variables in these seemingly disparate techniques can be mapped to each other.

816, TITLE: ROAD: Reality Oriented Adaptation for Semantic Segmentation of Urban Scenes
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_ROAD_Reality_Oriented_CVPR_2018_paper.html
AUTHORS: Yuhua Chen, Wen Li, Luc Van Gool
HIGHLIGHT: To this end, we propose a new reality oriented adaptation approach for urban scene semantic segmentation by learning from synthetic data.

817, TITLE: Eye In-Painting With Exemplar Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Dolhansky_Eye_In-Painting_With_CVPR_2018_paper.html
AUTHORS: Brian Dolhansky, Cristian Canton Ferrer

HIGHLIGHT: This paper introduces a novel approach to in-painting where the identity of the object to remove or change is preserved and accounted for at inference time: Exemplar GANs (ExGANs).

818, **TITLE:** ClcNet: Improving the Efficiency of Convolutional Neural Network Using Channel Local Convolutions
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_ClcNet_Improving_the_CVPR_2018_paper.html

AUTHORS: Dong-Qing Zhang

HIGHLIGHT: We suggest that these models can be considered as special cases of a generalized convolution operation, named channel local convolution (CLC), where an output channel is computed using a subset of the input channels.

819, **TITLE:** Towards Effective Low-Bitwidth Convolutional Neural Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Zhuang_Towards_Effective_Low-Bitwidth_CVPR_2018_paper.html

AUTHORS: Bohan Zhuang, Chunhua Shen, Mingkui Tan, Lingqiao Liu, Ian Reid

HIGHLIGHT: To mitigate this problem, we propose three simple-yet-effective approaches to improve the network training.

820, **TITLE:** Stochastic Downsampling for Cost-Adjustable Inference and Improved Regularization in Convolutional Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Kuen_Stochastic_Downsampling_for_CVPR_2018_paper.html

AUTHORS: Jason Kuen, Xiangfei Kong, Zhe Lin, Gang Wang, Jianxiong Yin, Simon See, Yap-Peng Tan

HIGHLIGHT: We propose a novel approach for cost-adjustable inference in CNNs - Stochastic Downsampling Point (SDPoint).

821, **TITLE:** Face Aging With Identity-Preserved Conditional Generative Adversarial Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Face_Aging_With_CVPR_2018_paper.html

AUTHORS: Zongwei Wang, Xu Tang, Weixin Luo, Shenghua Gao

HIGHLIGHT: Thus we propose an Identity-Preserved Conditional Generative Adversarial Networks (IPCAGNs) framework, in which a Conditional Generative Adversarial Networks module functions as generating a face that looks realistic and is with the target age, an identity-preserved module preserves the identity information and an age classifier forces the generated face with the target age.

822, **TITLE:** Unsupervised Cross-Dataset Person Re-Identification by Transfer Learning of Spatial-Temporal Patterns

http://openaccess.thecvf.com/content_cvpr_2018/html/Lv_Unsupervised_Cross-Dataset_Person_CVPR_2018_paper.html

AUTHORS: Jianming Lv, Weihang Chen, Qing Li, Can Yang

HIGHLIGHT: To address this challenge, we propose an unsupervised incremental learning algorithm, TFusion, which is aided by the transfer learning of the pedestrians' spatio-temporal patterns in the target domain.

823, **TITLE:** Feature Quantization for Defending Against Distortion of Images

http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Feature_Quantization_for_CVPR_2018_paper.html

AUTHORS: Zhun Sun, Mete Ozay, Yan Zhang, Xing Liu, Takayuki Okatani

HIGHLIGHT: In this work, we address the problem of improving robustness of convolutional neural networks (CNNs) to image distortion.

824, **TITLE:** Tagging Like Humans: Diverse and Distinct Image Annotation

http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Tagging_Like_Humans_CVPR_2018_paper.html

AUTHORS: Baoyuan Wu, Weidong Chen, Peng Sun, Wei Liu, Bernard Ghanem, Siwei Lyu

HIGHLIGHT: In this work we propose a new automatic image annotation model, dubbed diverse and distinct image annotation (D2IA).

In D2IA, we generate a relevant and distinct tag subset, in which the tags are relevant to the image contents and semantically distinct to each other, using sequential sampling from a determinantal point process (DPP) model.

825, **TITLE:** Re-Weighted Adversarial Adaptation Network for Unsupervised Domain Adaptation

http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Re-Weighted_Adversarial_Adaptation_CVPR_2018_paper.html

AUTHORS: Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, Kevin Chetty

HIGHLIGHT: In this paper, we propose the Re-weighted Adversarial Adaptation Network (RAAN) to reduce the feature distribution divergence and adapt the classifier when domain discrepancies are disparate.

826, **TITLE:** Inferring Semantic Layout for Hierarchical Text-to-Image Synthesis

http://openaccess.thecvf.com/content_cvpr_2018/html/Hong_Inferring_Semantic_Layout_CVPR_2018_paper.html

AUTHORS: Seunghoon Hong, Dingdong Yang, Jongwook Choi, Honglak Lee

HIGHLIGHT: We propose a novel hierarchical approach for text-to-image synthesis by inferring semantic layout.

- 827, TITLE: Regularizing RNNs for Caption Generation by Reconstructing the Past With the Present
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Regularizing_RNNs_for_CVPR_2018_paper.html
AUTHORS: Xinpeng Chen, Lin Ma, Wenhao Jiang, Jian Yao, Wei Liu
HIGHLIGHT: In this paper, we propose a novel architecture, namely Auto-Reconstructor Network (ARNet), which, coupling with the conventional encoder-decoder framework, works in an end-to-end fashion to generate captions.
- 828, TITLE: Unsupervised Domain Adaptation With Similarity Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Pinheiro_Unsupervised_Domain_Adaptation_CVPR_2018_paper.html
AUTHORS: Pedro O. Pinheiro
HIGHLIGHT: In this paper, we propose a different way to do the classification, using similarity learning.
- 829, TITLE: Learning Deep Sketch Abstraction
http://openaccess.thecvf.com/content_cvpr_2018/html/Muhammad_Learning_Deep_Sketch_CVPR_2018_paper.html
AUTHORS: Umar Riaz Muhammad, Yongxin Yang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales
HIGHLIGHT: In this work, we propose the first stroke-level sketch abstraction model based on the insight of sketch abstraction as a process of trading off between the recognizability of a sketch and the number of strokes used to draw it.
- 830, TITLE: Matching Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Mattyus_Matching_Adversarial_Networks_CVPR_2018_paper.html
AUTHORS: Gellért Mátyus, Raquel Urtasun
HIGHLIGHT: To overcome this, we propose to replace the discriminator with a matching network taking into account both the ground truth outputs as well as the generated examples.
- 831, TITLE: SoS-RSC: A Sum-of-Squares Polynomial Approach to Robustifying Subspace Clustering Algorithms
http://openaccess.thecvf.com/content_cvpr_2018/html/Sznaier_SoS-RSC_A_Sum-of-Squares_CVPR_2018_paper.html
AUTHORS: Mario Sznaier, Octavia Camps
HIGHLIGHT: To circumvent these difficulties, in this paper we propose an outlier removal algorithm based on evaluating a suitable sum-of-squares polynomial, computed directly from the data.
- 832, TITLE: Resource Aware Person Re-Identification Across Multiple Resolutions
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Resource_Aware_Person_CVPR_2018_paper.html
AUTHORS: Yan Wang, Lequn Wang, Yurong You, Xu Zou, Vincent Chen, Serena Li, Gao Huang, Bharath Hariharan, Kilian Q. Weinberger
HIGHLIGHT: To remedy this, we present a new person re-ID model that combines effective embeddings built on multiple convolutional network layers, trained with deep-supervision.
- 833, TITLE: Learning and Using the Arrow of Time
http://openaccess.thecvf.com/content_cvpr_2018/html/Wei_Learning_and_Using_CVPR_2018_paper.html
AUTHORS: Donglai Wei, Joseph J. Lim, Andrew Zisserman, William T. Freeman
HIGHLIGHT: To this end, we build three large-scale video datasets and apply a learning-based approach to these tasks.
- 834, TITLE: Neural Style Transfer via Meta Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Neural_Style_Transfer_CVPR_2018_paper.html
AUTHORS: Falong Shen, Shuicheng Yan, Gang Zeng
HIGHLIGHT: In this paper we propose a novel method to generate the specified network parameters through one feed-forward propagation in the meta networks for neural style transfer.
- 835, TITLE: People, Penguins and Petri Dishes: Adapting Object Counting Models to New Visual Domains and Object Types Without Forgetting
http://openaccess.thecvf.com/content_cvpr_2018/html/Marsden_People_Penguins_and_CVPR_2018_paper.html
AUTHORS: Mark Marsden, Kevin McGuinness, Suzanne Little, Ciara E. Keogh, Noel E. O'Connor
HIGHLIGHT: In this paper we propose a technique to adapt a convolutional neural network (CNN) based object counter to additional visual domains and object types while still preserving the original counting function.
- 836, TITLE: HydraNets: Specialized Dynamic Architectures for Efficient Inference
http://openaccess.thecvf.com/content_cvpr_2018/html/Mullapudi_HydraNets_Specialized_Dynamic_CVPR_2018_paper.html
AUTHORS: Ravi Teja Mullapudi, William R. Mark, Noam Shazeer, Kayvon Fatahalian
HIGHLIGHT: Specifically, we propose a network architecture template called HydraNet, which enables state-of-the-art architectures for image classification to be transformed into dynamic architectures which exploit conditional execution for efficient inference.

- 837, TITLE: SketchMate: Deep Hashing for Million-Scale Human Sketch Retrieval
http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_SketchMate_Deep_Hashing_CVPR_2018_paper.html
AUTHORS: Peng Xu, Yongye Huang, Tongtong Yuan, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales, Zhanyu Ma, Jun Guo
HIGHLIGHT: We propose a deep hashing framework for sketch retrieval that, for the first time, works on a multi-million scale human sketch dataset. Leveraging on this large dataset, we explore a few sketch-specific traits that were otherwise under-studied in prior literature.
Instead of following the conventional sketch recognition task, we introduce the novel problem of sketch hashing retrieval which is not only more challenging, but also offers a better testbed for large-scale sketch analysis, since: (i) more fine-grained sketch feature learning is required to accommodate the large variations in style and abstraction, and (ii) a compact binary code needs to be learned at the same time to enable efficient retrieval. Key to our network design is the embedding of unique characteristics of human sketch, where (i) a two-branch CNN-RNN architecture is adapted to explore the temporal ordering of strokes, and (ii) a novel hashing loss is specifically designed to accommodate both the temporal and abstract traits of sketches.
- 838, TITLE: From Source to Target and Back: Symmetric Bi-Directional Adaptive GAN
http://openaccess.thecvf.com/content_cvpr_2018/html/Russo_From_Source_to_CVPR_2018_paper.html
AUTHORS: Paolo Russo, Fabio M. Carlucci, Tatiana Tommasi, Barbara Caputo
HIGHLIGHT: In this paper we aim at getting the best of both worlds by introducing a symmetric mapping among domains.
- 839, TITLE: OLE: Orthogonal Low-Rank Embedding - A Plug and Play Geometric Loss for Deep Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Lezama_OLE_Orthogonal_Low-Rank_CVPR_2018_paper.html
AUTHORS: José Lezama, Qiang Qiu, Pablo Musé, Guillermo Sapiro
HIGHLIGHT: In this paper, we propose a plug-and-play loss term for deep networks that explicitly reduces intra-class variance and enforces inter-class margin simultaneously, in a simple and elegant geometric manner.
- 840, TITLE: Efficient Parametrization of Multi-Domain Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Rebuffi_Efficient_Parametrization_of_CVPR_2018_paper.html
AUTHORS: Sylvestre-Alvise Rebuffi, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: To overcome this limitation, in this paper we propose to consider instead universal parametric families of neural networks, which still contain specialized problem-specific models, but that differ only by a small number of parameters.
- 841, TITLE: Deep Density Clustering of Unconstrained Faces
http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_Deep_Density_Clustering_CVPR_2018_paper.html
AUTHORS: Wei-An Lin, Jun-Cheng Chen, Carlos D. Castillo, Rama Chellappa
HIGHLIGHT: In this paper, we consider the problem of grouping a collection of unconstrained face images in which the number of subjects is not known.
- 842, TITLE: Geometric Multi-Model Fitting With a Convex Relaxation Algorithm
http://openaccess.thecvf.com/content_cvpr_2018/html/Amayo_Geometric_Multi-Model_Fitting_CVPR_2018_paper.html
AUTHORS: Paul Amayo, Pedro Piniés, Lina M. Paz, Paul Newman
HIGHLIGHT: We propose a novel method for fitting multiple geometric models to multi-structural data via convex relaxation.
- 843, TITLE: Fast and Robust Estimation for Unit-Norm Constrained Linear Fitting Problems
http://openaccess.thecvf.com/content_cvpr_2018/html/Ikami_Fast_and_Robust_CVPR_2018_paper.html
AUTHORS: Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa
HIGHLIGHT: We overcome this problem by developing a novel objective function and its optimization, named iteratively reweighted eigenvalues minimization (IREM).
- 844, TITLE: Importance Weighted Adversarial Nets for Partial Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Importance_Weighted_Adversarial_CVPR_2018_paper.html
AUTHORS: Jing Zhang, Zewei Ding, Wanqing Li, Philip Ogunbona
HIGHLIGHT: This paper proposes an importance weighted adversarial nets-based method for unsupervised domain adaptation, specific for partial domain adaptation where the target domain has less number of classes compared to the source domain.
- 845, TITLE: Efficient Subpixel Refinement With Symbolic Linear Predictors
http://openaccess.thecvf.com/content_cvpr_2018/html/Lui_Efficient_Subpixel_Refinement_CVPR_2018_paper.html
AUTHORS: Vincent Lui, Jonathon Geeves, Winston Yui, Tom Drummond
HIGHLIGHT: We present an efficient subpixel refinement method using a learning-based approach called Linear Predictors.
- 846, TITLE: Scale-Recurrent Network for Deep Image Deblurring

- http://openaccess.thecvf.com/content_cvpr_2018/html/Tao_Scale-Recurrent_Network_for_CVPR_2018_paper.html
AUTHORS: Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, Jiaya Jia
HIGHLIGHT: In this paper, we investigate this strategy and propose a Scale-recurrent Network (SRN-DeblurNet) for this deblurring task.
- 847, TITLE: DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Kupyn_DeblurGAN_Blind_Motion_CVPR_2018_paper.html
AUTHORS: Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, Jiří Matas
HIGHLIGHT: We present DeblurGAN, an end-to-end learned method for motion deblurring.
- 848, TITLE: A2-RL: Aesthetics Aware Reinforcement Learning for Image Cropping
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_A2-RL_Aesthetics_Aware_CVPR_2018_paper.html
AUTHORS: Debang Li, Huikai Wu, Junge Zhang, Kaiqi Huang
HIGHLIGHT: Particularly, the proposed method develops an aesthetics aware reward function which especially benefits image cropping.
- 849, TITLE: Single Image Dehazing via Conditional Generative Adversarial Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Single_Image_DeHazing_CVPR_2018_paper.html
AUTHORS: Runde Li, Jinshan Pan, Zechao Li, Jinhui Tang
HIGHLIGHT: In this paper, we present an algorithm to directly restore a clear image from a hazy image.
- 850, TITLE: On the Duality Between Retinex and Image Dehazing
http://openaccess.thecvf.com/content_cvpr_2018/html/Galdran_On_the_Duality_CVPR_2018_paper.html
AUTHORS: Adrian Galdran, Aitor Alvarez-Gila, Alessandro Bria, Javier Vazquez-Corral, Marcelo Bertalmío
HIGHLIGHT: In this paper, we give theoretical proof that Retinex on inverted intensities is a solution to the image dehazing problem.
- 851, TITLE: Arbitrary Style Transfer With Deep Feature Reshuffle
http://openaccess.thecvf.com/content_cvpr_2018/html/Gu_Arbitrary_Style_Transfer_CVPR_2018_paper.html
AUTHORS: Shuyang Gu, Congliang Chen, Jing Liao, Lu Yuan
HIGHLIGHT: This paper introduces a novel method by reshuffling deep features (i.e., permuting the spacial locations of a feature map) of the style image for arbitrary style transfer.
- 852, TITLE: Nonlocal Low-Rank Tensor Factor Analysis for Image Restoration
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Nonlocal_Low-Rank_Tensor_CVPR_2018_paper.html
AUTHORS: Xinyuan Zhang, Xin Yuan, Lawrence Carin
HIGHLIGHT: We propose a new method that employs low-rank tensor factor analysis for tensors generated by grouped image patches.
- 853, TITLE: Avatar-Net: Multi-Scale Zero-Shot Style Transfer by Feature Decoration
http://openaccess.thecvf.com/content_cvpr_2018/html/Sheng_Avatar-Net_Multi-Scale_Zero-Shot_CVPR_2018_paper.html
AUTHORS: Lu Sheng, Ziyi Lin, Jing Shao, Xiaogang Wang
HIGHLIGHT: In this paper, we resolve this dilemma and propose an efficient yet effective Avatar-Net that enables visually plausible multi-scale transfer for arbitrary style.
- 854, TITLE: Missing Slice Recovery for Tensors Using a Low-Rank Model in Embedded Space
http://openaccess.thecvf.com/content_cvpr_2018/html/Yokota_Missing_Slice_Recovery_CVPR_2018_paper.html
AUTHORS: Tatsuya Yokota, Burak Erem, Seyhmus Guler, Simon K. Warfield, Hidekata Hontani
HIGHLIGHT: In this study, we consider a low-rank model in an embedded space of a tensor.
- 855, TITLE: Deep Semantic Face Deblurring
http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_Deep_Semantic_Face_CVPR_2018_paper.html
AUTHORS: Ziyi Shen, Wei-Sheng Lai, Tingfa Xu, Jan Kautz, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we present an effective and efficient face deblurring algorithm by exploiting semantic cues via deep convolutional neural networks (CNNs).
- 856, TITLE: GraphBit: Bitwise Interaction Mining via Deep Reinforcement Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Duan_GraphBit_Bitwise_Interaction_CVPR_2018_paper.html
AUTHORS: Yueqi Duan, Ziwei Wang, Jiwen Lu, Xudong Lin, Jie Zhou

HIGHLIGHT: In this paper, we propose a GraphBit method to learn deep binary descriptors in a directed acyclic graph unsupervisedly, representing bitwise interactions as edges between the nodes of bits.

857, **TITLE:** Recurrent Saliency Transformation Network: Incorporating Multi-Stage Visual Cues for Small Organ Segmentation

http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_Recurrent_Saliency_Transformation_CVPR_2018_paper.html

AUTHORS: Qihang Yu, Lingxi Xie, Yan Wang, Yuyin Zhou, Elliot K. Fishman, Alan L. Yuille

HIGHLIGHT: We aim at segmenting small organs (e.g., the pancreas) from abdominal CT scans.

858, **TITLE:** Thoracic Disease Identification and Localization With Limited Supervision

http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Thoracic_Disease_Identification_CVPR_2018_paper.html

AUTHORS: Zhe Li, Chong Wang, Mei Han, Yuan Xue, Wei Wei, Li-Jia Li, Li Fei-Fei

HIGHLIGHT: To address this challenge, we present a unified approach that simultaneously performs disease identification and localization through the same underlying model for all images. We demonstrate that our approach can effectively leverage both class information as well as limited location annotation, and significantly outperforms the comparative reference baseline in both classification and localization tasks.

859, **TITLE:** Quantization of Fully Convolutional Networks for Accurate Biomedical Image Segmentation

http://openaccess.thecvf.com/content_cvpr_2018/html/Xu_Quantization_of_Fully_CVPR_2018_paper.html

AUTHORS: Xiaowei Xu, Qing Lu, Lin Yang, Sharon Hu, Danny Chen, Yu Hu, Yiyu Shi

HIGHLIGHT: In this paper, we apply quantization techniques to FCNs for accurate biomedical image segmentation.

860, **TITLE:** Visual Feature Attribution Using Wasserstein GANs

http://openaccess.thecvf.com/content_cvpr_2018/html/Baumgartner_Visual_Feature_Attribution_CVPR_2018_paper.html

AUTHORS: Christian F. Baumgartner, Lisa M. Koch, Kerem Can Tezcan, Jia Xi Ang, Ender Konukoglu

HIGHLIGHT: In this paper, we discuss a limitation of these approaches which may lead to only a subset of the category specific features being detected.

861, **TITLE:** Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies

http://openaccess.thecvf.com/content_cvpr_2018/html/Joo_Total_Capture_A_CVPR_2018_paper.html

AUTHORS: Hanbyul Joo, Tomas Simon, Yaser Sheikh

HIGHLIGHT: We present a unified deformation model for the markerless capture of multiple scales of human movement, including facial expressions, body motion, and hand gestures.

862, **TITLE:** Augmented Skeleton Space Transfer for Depth-Based Hand Pose Estimation

http://openaccess.thecvf.com/content_cvpr_2018/html/Baek_Augmented_Skeleton_Space_CVPR_2018_paper.html

AUTHORS: Seungryul Baek, Kwang In Kim, Tae-Kyun Kim

HIGHLIGHT: Crucial to the success of training a depth-based 3D hand pose estimator (HPE) is the availability of comprehensive datasets covering diverse camera perspectives, shapes, and pose variations.

863, **TITLE:** Synthesizing Images of Humans in Unseen Poses

http://openaccess.thecvf.com/content_cvpr_2018/html/Balakrishnan_Synthesizing_Images_of_CVPR_2018_paper.html

AUTHORS: Guha Balakrishnan, Amy Zhao, Adrian V. Dalca, Frédo Durand, John Guttag

HIGHLIGHT: We present a modular generative neural network that synthesizes unseen poses using training pairs of images and poses taken from human action videos.

864, **TITLE:** SSNet: Scale Selection Network for Online 3D Action Prediction

http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_SSNet_Scale_Selection_CVPR_2018_paper.html

AUTHORS: Jun Liu, Amir Shahroudy, Gang Wang, Ling-Yu Duan, Alex C. Kot

HIGHLIGHT: In this paper, we focus on online action prediction in streaming 3D skeleton sequences.

865, **TITLE:** Detecting and Recognizing Human-Object Interactions

http://openaccess.thecvf.com/content_cvpr_2018/html/Gkioxari_Detecting_and_Recognizing_CVPR_2018_paper.html

AUTHORS: Georgia Gkioxari, Ross Girshick, Piotr Dollár, Kaiming He

HIGHLIGHT: In this paper, we address the task of detecting (human, verb, object) triplets in challenging everyday photos.

866, **TITLE:** Unsupervised Learning and Segmentation of Complex Activities From Video

http://openaccess.thecvf.com/content_cvpr_2018/html/Sener_Unsupervised_Learning_and_CVPR_2018_paper.html

AUTHORS: Fadime Sener, Angela Yao

HIGHLIGHT: This paper presents a new method for unsupervised segmentation of complex activities from video into multiple steps, or sub-activities, without any textual input.

867, **TITLE:** Unsupervised Training for 3D Morphable Model Regression
http://openaccess.thecvf.com/content_cvpr_2018/html/Genova_Unsupervised_Training_for_CVPR_2018_paper.html
AUTHORS: Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, William T. Freeman
HIGHLIGHT: To make training from features feasible and avoid network fooling effects, we introduce three objectives: a batch distribution loss that encourages the output distribution to match the distribution of the morphable model, a loopback loss that ensures the network can correctly reinterpret its own output, and a multi-view identity loss that compares the features of the predicted 3D face and the input photograph from multiple viewing angles.

868, **TITLE:** Video Based Reconstruction of 3D People Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Alldieck_Video_Based_Reconstruction_CVPR_2018_paper.html
AUTHORS: Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, Gerard Pons-Moll
HIGHLIGHT: This paper describes how to obtain accurate 3D body models and texture of arbitrary people from a single, monocular video in which a person is moving.

869, **TITLE:** Pose-Guided Photorealistic Face Rotation
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Pose-Guided_Photorealistic_Face_CVPR_2018_paper.html
AUTHORS: Yibo Hu, Xiang Wu, Bing Yu, Ran He, Zhenan Sun
HIGHLIGHT: We propose a novel Couple-Agent Pose-Guided Generative Adversarial Network (CAPG-GAN) to generate both neutral and profile head pose face images.

870, **TITLE:** Mesoscopic Facial Geometry Inference Using Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Huynh_Mesoscopic_Facial_Geometry_CVPR_2018_paper.html
AUTHORS: Loc Huynh, Weikai Chen, Shunsuke Saito, Jun Xing, Koki Nagano, Andrew Jones, Paul Debevec, Hao Li
HIGHLIGHT: We present a learning-based approach for synthesizing facial geometry at medium and fine scales from diffusely-lit facial texture maps.

871, **TITLE:** Hand PointNet: 3D Hand Pose Estimation Using Point Sets
http://openaccess.thecvf.com/content_cvpr_2018/html/Ge_Hand_PointNet_3D_CVPR_2018_paper.html
AUTHORS: Lihao Ge, Yujun Cai, Junwu Weng, Junsong Yuan
HIGHLIGHT: Different from existing CNN-based hand pose estimation methods that take either 2D images or 3D volumes as the input, our proposed Hand PointNet directly processes the 3D point cloud that models the visible surface of the hand for pose regression.

872, **TITLE:** Seeing Voices and Hearing Faces: Cross-Modal Biometric Matching
http://openaccess.thecvf.com/content_cvpr_2018/html/Nagrani_Seeing_Voices_and_CVPR_2018_paper.html
AUTHORS: Arsha Nagrani, Samuel Albanie, Andrew Zisserman
HIGHLIGHT: In this paper we study this, and a number of related cross-modal tasks, aimed at answering the question: how much can we infer from the voice about the face and vice versa?

873, **TITLE:** Learning Monocular 3D Human Pose Estimation From Multi-View Images
http://openaccess.thecvf.com/content_cvpr_2018/html/Rhodin_Learning_Monocular_3D_CVPR_2018_paper.html
AUTHORS: Helge Rhodin, Jörg Spörrri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: In this paper, we propose to replace most of the annotations by the use of multiple views, at training time only.

874, **TITLE:** Separating Style and Content for Generalized Style Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Separating_Style_and_CVPR_2018_paper.html
AUTHORS: Yexun Zhang, Ya Zhang, Wenbin Cai
HIGHLIGHT: We here attempt to separate the representations for styles and contents, and propose a generalized style transfer network consisting of style encoder, content encoder, mixer and decoder.

875, **TITLE:** TextureGAN: Controlling Deep Image Synthesis With Texture Patches
http://openaccess.thecvf.com/content_cvpr_2018/html/Xian_TextureGAN_Controlling_Deep_CVPR_2018_paper.html
AUTHORS: Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, James Hays
HIGHLIGHT: In this paper, we investigate deep image synthesis guided by sketch, color, and texture.

876, **TITLE:** Connecting Pixels to Privacy and Utility: Automatic Redaction of Private Information in Images

http://openaccess.thecvf.com/content_cvpr_2018/html/Orekondy_Connecting_Pixels_to_CVPR_2018_paper.html

AUTHORS: Tribhuvanesh Orekondy, Mario Fritz, Bernt Schiele

HIGHLIGHT: We present the first model for automatic redaction of diverse private information.

877, TITLE: MapNet: An Allocentric Spatial Memory for Mapping Environments

http://openaccess.thecvf.com/content_cvpr_2018/html/Henriques_MapNet_An_Allocentric_CVPR_2018_paper.html

AUTHORS: João F. Henriques, Andrea Vedaldi

HIGHLIGHT: In this paper, we develop a differentiable module that satisfies such requirements, while being robust, efficient, and suitable for integration in end-to-end deep networks.

878, TITLE: Accurate and Diverse Sampling of Sequences Based on a “Best of Many” Sample Objective

http://openaccess.thecvf.com/content_cvpr_2018/html/Bhattacharyya_Accurate_and_Diverse_CVPR_2018_paper.html

AUTHORS: Apratim Bhattacharyya, Bernt Schiele, Mario Fritz

HIGHLIGHT: Our core contribution is a “Best of Many” sample objective that leads to more accurate and more diverse predictions that better capture the true variations in real-world sequence data.

879, TITLE: VirtualHome: Simulating Household Activities via Programs

http://openaccess.thecvf.com/content_cvpr_2018/html/Puig_VirtualHome_Simulating_Household_CVPR_2018_paper.html

AUTHORS: Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, Antonio Torralba

HIGHLIGHT: In this paper, we are interested in modeling complex activities that occur in a typical household.

880, TITLE: Generate to Adapt: Aligning Domains Using Generative Adversarial Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Sankaranarayanan_Generate_to_Adapt_CVPR_2018_paper.html

AUTHORS: Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, Rama Chellappa

HIGHLIGHT: In this work, we propose an approach that leverages unsupervised data to bring the source and target distributions closer in a learned joint feature space.

881, TITLE: Multi-Agent Diverse Generative Adversarial Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Ghosh_Multi-Agent_Diverse_Generative_CVPR_2018_paper.html

AUTHORS: Arnab Ghosh, Viveka Kulharia, Vinay P. Namboodiri, Philip H.S. Torr, Puneet K. Dokania

HIGHLIGHT: We propose MAD-GAN, an intuitive generalization to the Generative Adversarial Networks (GANs) and its conditional variants to address the well known problem of mode collapse.

882, TITLE: A PID Controller Approach for Stochastic Optimization of Deep Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/An_A_PID_Controller_CVPR_2018_paper.html

AUTHORS: Wangpeng An, Haoqian Wang, Qingyun Sun, Jun Xu, Qionghai Dai, Lei Zhang

HIGHLIGHT: Inspired by the prominent success of proportional-integral-derivative (PID) controller in automatic control, we propose a PID approach for accelerating deep network optimization.

883, TITLE: “Learning-Compression” Algorithms for Neural Net Pruning

http://openaccess.thecvf.com/content_cvpr_2018/html/Carreira-Perpinan_Learning-Compression_Algorithms_for_CVPR_2018_paper.html

AUTHORS: Miguel Á. Carreira-Perpiñán, Yerlan Idelbayev

HIGHLIGHT: We formulate pruning as an optimization problem of finding the weights that minimize the loss while satisfying a pruning cost condition.

884, TITLE: Large-Scale Distance Metric Learning With Uncertainty

http://openaccess.thecvf.com/content_cvpr_2018/html/Qian_Large-Scale_Distance_Metric_CVPR_2018_paper.html

AUTHORS: Qi Qian, Jiasheng Tang, Hao Li, Shenghuo Zhu, Rong Jin

HIGHLIGHT: In this work, we propose the margin preserving metric learning framework to learn the distance metric and latent examples simultaneously.

885, TITLE: Guide Me: Interacting With Deep Networks

http://openaccess.thecvf.com/content_cvpr_2018/html/Rupprecht_Guide_Me_Interacting_CVPR_2018_paper.html

AUTHORS: Christian Rupprecht, Iro Laina, Nassir Navab, Gregory D. Hager, Federico Tombari

HIGHLIGHT: In this paper, we explore methods to flexibly guide a trained convolutional neural network through user input to improve its performance during inference.

886, TITLE: Art of Singular Vectors and Universal Adversarial Perturbations

http://openaccess.thecvf.com/content_cvpr_2018/html/Khrulkov_Art_of_Singular_CVPR_2018_paper.html

AUTHORS: Valentin Khrulkov, Ivan Oseledets
HIGHLIGHT: In this work we propose a new algorithm for constructing such universal perturbations.

887, TITLE: Deflecting Adversarial Attacks With Pixel Deflection
http://openaccess.thecvf.com/content_cvpr_2018/html/Prakash_Deflecting_Adversarial_Attacks_CVPR_2018_paper.html
AUTHORS: Aaditya Prakash, Nick Moran, Solomon Garber, Antonella DiLillo, James Storer
HIGHLIGHT: We present an algorithm to process an image so that classification accuracy is significantly preserved in the presence of such adversarial manipulations.

888, TITLE: MovieGraphs: Towards Understanding Human-Centric Situations From Videos
http://openaccess.thecvf.com/content_cvpr_2018/html/Vicol_MovieGraphs_Towards_Understanding_CVPR_2018_paper.html
AUTHORS: Paul Vicol, Makarand Tapaswi, Lluís Castrejón, Sanja Fidler
HIGHLIGHT: We propose a method for querying videos and text with graphs, and show that: 1) our graphs contain rich and sufficient information to summarize and localize each scene; and 2) subgraphs allow us to describe situations at an abstract level and retrieve multiple semantically relevant situations.
Towards this goal, we introduce a novel dataset called MovieGraphs which provides detailed, graph-based annotations of social situations depicted in movie clips.

889, TITLE: SemStyle: Learning to Generate Stylised Image Captions Using Unaligned Text
http://openaccess.thecvf.com/content_cvpr_2018/html/Mathews_SemStyle_Learning_to_CVPR_2018_paper.html
AUTHORS: Alexander Mathews, Lexing Xie, Xuming He
HIGHLIGHT: We develop a model that learns to generate visually relevant styled captions from a large corpus of styled text without aligned images.

890, TITLE: Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions
http://openaccess.thecvf.com/content_cvpr_2018/html/Sattler_Benchmarking_6DOF_Outdoor_CVPR_2018_paper.html
AUTHORS: Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, Tomas Pajdla
HIGHLIGHT: In this paper, we introduce the first benchmark datasets specifically designed for analyzing the impact of such factors on visual localization.

891, TITLE: IVQA: Inverse Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_IVQA_Inverse_Visual_CVPR_2018_paper.html
AUTHORS: Feng Liu, Tao Xiang, Timothy M. Hospedales, Wankou Yang, Changyin Sun
HIGHLIGHT: We pose question generation as a multi-modal dynamic inference process and propose an iVQA model that can gradually adjust its focus of attention guided by both a partially generated question and the answer.
We propose the inverse problem of Visual question answering (iVQA), and explore its suitability as a benchmark for visuo-linguistic understanding.

892, TITLE: Unsupervised Person Image Synthesis in Arbitrary Poses
http://openaccess.thecvf.com/content_cvpr_2018/html/Pumarola_Unsupervised_Person_Image_CVPR_2018_paper.html
AUTHORS: Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, Francesc Moreno-Noguer
HIGHLIGHT: We present a novel approach for synthesizing photo-realistic images of people in arbitrary poses using generative adversarial learning.

893, TITLE: Learning Descriptor Networks for 3D Shape Synthesis and Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Xie_Learning_Descriptor_Networks_CVPR_2018_paper.html
AUTHORS: Jianwen Xie, Zilong Zheng, Ruiqi Gao, Wenguan Wang, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT: This paper proposes a 3D shape descriptor network, which is a deep convolutional energy-based model, for modeling volumetric shape patterns.

894, TITLE: Neural Kinematic Networks for Unsupervised Motion Retargetting
http://openaccess.thecvf.com/content_cvpr_2018/html/Villegas_Neural_Kinematic_Networks_CVPR_2018_paper.html
AUTHORS: Ruben Villegas, Jimei Yang, Duygu Ceylan, Honglak Lee
HIGHLIGHT: We propose a recurrent neural network architecture with a Forward Kinematics layer and cycle consistency based adversarial training objective for unsupervised motion retargetting.

895, TITLE: Group Consistent Similarity Learning via Deep CRF for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Group_Consistent_Similarity_CVPR_2018_paper.html
AUTHORS: Dapeng Chen, Dan Xu, Hongsheng Li, Nicu Sebe, Xiaogang Wang

HIGHLIGHT: In this paper, we incorporate constraints on large image groups by combining the CRF with deep neural networks.

896, **TITLE:** Learning Compositional Visual Concepts With Mutual Consistency
http://openaccess.thecvf.com/content_cvpr_2018/html/Gong_Learning_Compositional_Visual_CVPR_2018_paper.html
AUTHORS: Yunye Gong, Srikrishna Karanam, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, Peter C. Doerschuk
HIGHLIGHT: We present a novel answer in this paper based on cyclic consistency over multiple concepts, represented individually by generative adversarial networks (GANs).

897, **TITLE:** NestedNet: Learning Nested Sparse Structures in Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_NestedNet_Learning_Nested_CVPR_2018_paper.html
AUTHORS: Eunwoo Kim, Chanho Ahn, Songhwai Oh
HIGHLIGHT: In this work, we propose a novel deep learning framework, called a nested sparse network, which exploits an n-in-1-type nested structure in a neural network.

898, **TITLE:** Context Embedding Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Kim_Context_Embedding_Networks_CVPR_2018_paper.html
AUTHORS: Kun Ho Kim, Oisín Mac Aodha, Pietro Perona
HIGHLIGHT: To overcome these limitations we introduce Context Embedding Networks (CENs).

899, **TITLE:** Iterative Learning With Open-Set Noisy Labels
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Iterative_Learning_With_CVPR_2018_paper.html
AUTHORS: Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, Shu-Tao Xia
HIGHLIGHT: To address this problem, we propose a novel iterative learning framework for training CNNs on datasets with open-set noisy labels.

900, **TITLE:** Learning Transferable Architectures for Scalable Image Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zoph_Learning_Transferable_Architectures_CVPR_2018_paper.html
AUTHORS: Barret Zoph, Vijay Vasudevan, Jonathon Shlens, Quoc V. Le
HIGHLIGHT: In this paper, we study a method to learn the model architectures directly on the dataset of interest.

901, **TITLE:** SBNet: Sparse Blocks Network for Fast Inference
http://openaccess.thecvf.com/content_cvpr_2018/html/Ren_SBNet_Sparse_Blocks_CVPR_2018_paper.html
AUTHORS: Mengye Ren, Andrei Pokrovsky, Bin Yang, Raquel Urtasun
HIGHLIGHT: In this work, we leverage the sparsity structure of computation masks and propose a novel tiling-based sparse convolution algorithm.

902, **TITLE:** Language-Based Image Editing With Recurrent Attentive Models
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Language-Based_Image_Editing_CVPR_2018_paper.html
AUTHORS: Jianbo Chen, Yelong Shen, Jianfeng Gao, Jingjing Liu, Xiaodong Liu
HIGHLIGHT: We propose a generic modeling framework for two sub-tasks of LBIE: language-based image segmentation and image colorization.
First, we introduce a synthetic dataset, called CoSaL, to evaluate the end-to-end performance of our LBIE system.

903, **TITLE:** Net2Vec: Quantifying and Explaining How Concepts Are Encoded by Filters in Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Fong_Net2Vec_Quantifying_and_CVPR_2018_paper.html
AUTHORS: Ruth Fong, Andrea Vedaldi
HIGHLIGHT: In order to investigate this idea while enabling systematic visualization and quantification of multiple filter responses, we introduce the Net2Vec framework, in which semantic concepts are mapped to vectorial embeddings based on corresponding filter responses.

904, **TITLE:** End-to-End Dense Video Captioning With Masked Transformer
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_End-to-End_Dense_Video_CVPR_2018_paper.html
AUTHORS: Luowei Zhou, Yingbo Zhou, Jason J. Corso, Richard Socher, Caiming Xiong
HIGHLIGHT: To address this problem, we propose an end-to-end transformer model for dense video captioning.

905, **TITLE:** A Neural Multi-Sequence Alignment TeCHnique (NeuMATCH)
http://openaccess.thecvf.com/content_cvpr_2018/html/Dogan_A_Neural_Multi-Sequence_CVPR_2018_paper.html
AUTHORS: Pelin Dogan, Boyang Li, Leonid Sigal, Markus Gross

HIGHLIGHT: In this paper, we propose an end-to-end neural architecture where alignment actions are implemented as moving data between stacks of Long Short-term Memory (LSTM) blocks.

906, **TITLE:** Path Aggregation Network for Instance Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Path_Aggregation_Network_CVPR_2018_paper.html
AUTHORS: Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, Jiaya Jia
HIGHLIGHT: In this paper, we propose Path Aggregation Network (PANet) aiming at boosting information flow in proposal-based instance segmentation framework.

907, **TITLE:** The iNaturalist Species Classification and Detection Dataset
http://openaccess.thecvf.com/content_cvpr_2018/html/Van_Horn_The_iNaturalist_Species_CVPR_2018_paper.html
AUTHORS: Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, Serge Belongie
HIGHLIGHT: To encourage further progress in challenging real world conditions we present the iNaturalist species classification and detection dataset, consisting of 859,000 images from over 5,000 different species of plants and animals.

908, **TITLE:** Multimodal Explanations: Justifying Decisions and Pointing to the Evidence
http://openaccess.thecvf.com/content_cvpr_2018/html/Park_Multimodal_Explanations_Justifying_CVPR_2018_paper.html
AUTHORS: Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, Marcus Rohrbach
HIGHLIGHT: We propose a multimodal approach to explanation, and argue that the two modalities provide complementary explanatory strengths. We collect two new datasets to define and evaluate this task, and propose a novel model which can provide joint textual rationale generation and attention visualization.

909, **TITLE:** StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation
http://openaccess.thecvf.com/content_cvpr_2018/html/Choi_StarGAN_Unified_Generative_CVPR_2018_paper.html
AUTHORS: Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, Jaegul Choo
HIGHLIGHT: To address this limitation, we propose StarGAN, a novel and scalable approach that can perform image-to-image translations for multiple domains using only a single model.

910, **TITLE:** High-Resolution Image Synthesis and Semantic Manipulation With Conditional GANs
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_High-Resolution_Image_Synthesis_CVPR_2018_paper.html
AUTHORS: Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, Bryan Catanzaro
HIGHLIGHT: We present a new method for synthesizing high-resolution photo-realistic images from semantic label maps using conditional generative adversarial networks (conditional GANs).

911, **TITLE:** Semi-Parametric Image Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Semi-Parametric_Image_Synthesis_CVPR_2018_paper.html
AUTHORS: Xiaojuan Qi, Qifeng Chen, Jiaya Jia, Vladlen Koltun
HIGHLIGHT: We present a semi-parametric approach to photographic image synthesis from semantic layouts.

912, **TITLE:** BlockDrop: Dynamic Inference Paths in Residual Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_BlockDrop_Dynamic_Inference_CVPR_2018_paper.html
AUTHORS: Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S. Davis, Kristen Grauman, Rogerio Feris
HIGHLIGHT: We introduce BlockDrop, an approach that learns to dynamically choose which layers of a deep network to execute during inference so as to best reduce total computation without degrading prediction accuracy.

913, **TITLE:** Interpretable Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Interpretable_Convolutional_Neural_CVPR_2018_paper.html
AUTHORS: Quanshi Zhang, Ying Nian Wu, Song-Chun Zhu
HIGHLIGHT: This paper proposes a method to modify a traditional convolutional neural network (CNN) into an interpretable CNN, in order to clarify knowledge representations in high conv-layers of the CNN.

914, **TITLE:** Deep Cross-Media Knowledge Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Huang_Deep_Cross-Media_Knowledge_CVPR_2018_paper.html
AUTHORS: Xin Huang, Yuxin Peng
HIGHLIGHT: For achieving the goal, this paper proposes deep cross-media knowledge transfer (DCKT) approach, which transfers knowledge from a large-scale cross-media dataset to promote the model training on another small-scale cross-media dataset.

- 915, TITLE: Interleaved Structured Sparse Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Xie_Interleaved_Structured_Sparse_CVPR_2018_paper.html
AUTHORS: Guotian Xie, Jingdong Wang, Ting Zhang, Jianhuang Lai, Richang Hong, Guo-Jun Qi
HIGHLIGHT: In this paper, we study the problem of designing efficient convolutional neural network architectures with the interest in eliminating the redundancy in convolution kernels.
- 916, TITLE: A Variational U-Net for Conditional Appearance and Shape Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Esser_A_Variational_U-Net_CVPR_2018_paper.html
AUTHORS: Patrick Esser, Ekaterina Sutter, Björn Ommer
HIGHLIGHT: We present a conditional U-Net for shape-guided image generation, conditioned on the output of a variational autoencoder for appearance.
- 917, TITLE: Detach and Adapt: Learning Cross-Domain Disentangled Deep Representation
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Detach_and_Adapt_CVPR_2018_paper.html
AUTHORS: Yen-Cheng Liu, Yu-Ying Yeh, Tzu-Chien Fu, Sheng-De Wang, Wei-Chen Chiu, Yu-Chiang Frank Wang
HIGHLIGHT: To address this problem, we propose a novel deep learning model of Cross-Domain Representation Disentangler (CDRD).
- 918, TITLE: Learning Deep Structured Active Contours End-to-End
http://openaccess.thecvf.com/content_cvpr_2018/html/Marcos_Learning_Deep_Structured_CVPR_2018_paper.html
AUTHORS: Diego Marcos, Devis Tuia, Benjamin Kellenberger, Lisa Zhang, Min Bai, Renjie Liao, Raquel Urtasun
HIGHLIGHT: To this end, we present Deep Structured Active Contours (DSAC), a novel framework that integrates priors and constraints into the segmentation process, such as continuous boundaries, smooth edges, and sharp corners.
- 919, TITLE: Deep Learning Under Privileged Information Using Heteroscedastic Dropout
http://openaccess.thecvf.com/content_cvpr_2018/html/Lambert_Deep_Learning_Under_CVPR_2018_paper.html
AUTHORS: John Lambert, Ozan Sener, Silvio Savarese
HIGHLIGHT: We propose a new LUPI algorithm specifically designed for Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs).
- 920, TITLE: Smooth Neighbors on Teacher Graphs for Semi-Supervised Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Luo_Smooth_Neighbors_on_CVPR_2018_paper.html
AUTHORS: Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, Bo Zhang
HIGHLIGHT: In this paper, we propose a novel method, called Smooth Neighbors on Teacher Graphs (SNTG).
- 921, TITLE: Interpret Neural Networks by Identifying Critical Data Routing Paths
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Interpret_Neural_Networks_CVPR_2018_paper.html
AUTHORS: Yulong Wang, Hang Su, Bo Zhang, Xiaolin Hu
HIGHLIGHT: Based on the discoveries, we propose an adversarial sample detection algorithm by learning a classifier to discriminate whether the critical data routing paths are from real or adversarial samples.
- 922, TITLE: Deep Spatio-Temporal Random Fields for Efficient Video Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Chandra_Deep_Spatio-Temporal_Random_CVPR_2018_paper.html
AUTHORS: Siddhartha Chandra, Camille Couprie, Iasonas Kokkinos
HIGHLIGHT: In this work we introduce a time- and memory-efficient method for structured prediction that couples neuron decisions across both space at time.
- 923, TITLE: Customized Image Narrative Generation via Interactive Visual Question Generation and Answering
http://openaccess.thecvf.com/content_cvpr_2018/html/Shin_Customized_Image_Narrative_CVPR_2018_paper.html
AUTHORS: Andrew Shin, Yoshitaka Ushiku, Tatsuya Harada
HIGHLIGHT: In this paper, we propose a customized image narrative generation task, in which the users are interactively engaged in the generation process by providing answers to the questions.
- 924, TITLE: PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_PWC-Net_CNNs_for_CVPR_2018_paper.html
AUTHORS: Deqing Sun, Xiaodong Yang, Ming-Yu Liu, Jan Kautz
HIGHLIGHT: We present a compact but effective CNN model for optical flow, called PWC-Net.
- 925, TITLE: Revisiting Deep Intrinsic Image Decompositions

- http://openaccess.thecvf.com/content_cvpr_2018/html/Fan_Revisiting_Deep_Intrinsic_CVPR_2018_paper.html
AUTHORS: Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, David Wipf
HIGHLIGHT: In contrast to many previous learning-based approaches, which are often tailored to the structure of a particular dataset (and may not work well on others), we adopt core network structures that universally reflect loose prior knowledge regarding the intrinsic image formation process and can be largely shared across datasets.
- 926, TITLE: Multi-Cell Detection and Classification Using a Generative Convolutional Model
http://openaccess.thecvf.com/content_cvpr_2018/html/Yellin_Multi-Cell_Detection_and_CVPR_2018_paper.html
AUTHORS: Florence Yellin, Benjamin D. Haeffele, Sophie Roth, René Vidal
HIGHLIGHT: This paper proposes a new approach to detecting, counting and classifying white blood cell populations in holographic images, which capitalizes on the fact that the variability in a mixture of blood cells is constrained by physiology.
- 927, TITLE: Learning Spatial-Aware Regressions for Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2018/html/Sun_Learning_Spatial-Aware_Regressions_CVPR_2018_paper.html
AUTHORS: Chong Sun, Dong Wang, Huchuan Lu, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we analyze the spatial information of deep features, and propose two complementary regressions for robust visual tracking.
- 928, TITLE: High Performance Visual Tracking With Siamese Region Proposal Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_High_Performance_Visual_CVPR_2018_paper.html
AUTHORS: Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, Xiaolin Hu
HIGHLIGHT: In this paper, we propose the Siamese region proposal network (Siamese-RPN) which is end-to-end trained off-line with large-scale image pairs.
- 929, TITLE: LiteFlowNet: A Lightweight Convolutional Neural Network for Optical Flow Estimation
http://openaccess.thecvf.com/content_cvpr_2018/html/Hui_LiteFlowNet_A_Lightweight_CVPR_2018_paper.html
AUTHORS: Tak-Wai Hui, Xiaoou Tang, Chen Change Loy
HIGHLIGHT: In this paper we present an alternative network that attains performance on par with FlowNet2 on the challenging Sintel final pass and KITTI benchmarks, while being 30 times smaller in the model size and 1.36 times faster in the running speed.
- 930, TITLE: VITAL: Visual Tracking via Adversarial Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Song_VITAL_Visual_Tracking_CVPR_2018_paper.html
AUTHORS: Yibing Song, Chao Ma, Xiaohe Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson W.H. Lau, Ming-Hsuan Yang
HIGHLIGHT: This paper presents the VITAL algorithm to address these two problems via adversarial learning.
- 931, TITLE: Super SloMo: High Quality Estimation of Multiple Intermediate Frames for Video Interpolation
http://openaccess.thecvf.com/content_cvpr_2018/html/Jiang_Super_SloMo_High_CVPR_2018_paper.html
AUTHORS: Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, Jan Kautz
HIGHLIGHT: We start by computing bi-directional optical flow between the input images using a U-Net architecture.
- 932, TITLE: Real-World Repetition Estimation by Div, Grad and Curl
http://openaccess.thecvf.com/content_cvpr_2018/html/Runia_Real-World_Repetition_Estimation_CVPR_2018_paper.html
AUTHORS: Tom F. H. Runia, Cees G. M. Snoek, Arnold W. M. Smeulders
HIGHLIGHT: We consider the problem of estimating repetition in video, such as performing push-ups, cutting a melon or playing violin.
For experiments, we introduce the new QUVA Repetition dataset, reflecting reality by including non-static and non-stationary videos.
- 933, TITLE: Recurrent Pixel Embedding for Instance Grouping
http://openaccess.thecvf.com/content_cvpr_2018/html/Kong_Recurrent_Pixel_Embedding_CVPR_2018_paper.html
AUTHORS: Shu Kong, Charless C. Fowlkes
HIGHLIGHT: We introduce a differentiable, end-to-end trainable framework for solving pixel-level grouping problems such as instance segmentation consisting of two novel components.
- 934, TITLE: Deep Unsupervised Saliency Detection: A Multiple Noisy Labeling Perspective
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Deep_Unsupervised_Saliency_CVPR_2018_paper.html
AUTHORS: Jing Zhang, Tong Zhang, Yuchao Dai, Mehrtash Harandi, Richard Hartley
HIGHLIGHT: To this end, we present a novel perspective to unsupervised saliency detection through learning from multiple noisy labeling generated by "weak" and "noisy" unsupervised handcrafted saliency methods.

935, TITLE: Learning Intrinsic Image Decomposition From Watching the World
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Learning_Intrinsic_Image_CVPR_2018_paper.html
AUTHORS: Zhengqi Li, Noah Snavely
HIGHLIGHT: In this paper, we explore a different approach to learning intrinsic images: observing image sequences over time depicting the same scene under changing illumination, and learning single-view decompositions that are consistent with these changes.

936, TITLE: TieNet: Text-Image Embedding Network for Common Thorax Disease Classification and Reporting in Chest X-Rays
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_TieNet_Text-Image_Embedding_CVPR_2018_paper.html
AUTHORS: Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Ronald M. Summers
HIGHLIGHT: In this paper, we show the clinical free-text radiological reports can be utilized as a priori knowledge for tackling these two key problems.

937, TITLE: Generating Synthetic X-Ray Images of a Person From the Surface Geometry
http://openaccess.thecvf.com/content_cvpr_2018/html/Teixeira_Generating_Synthetic_X-Ray_CVPR_2018_paper.html
AUTHORS: Brian Teixeira, Vivek Singh, Terrence Chen, Kai Ma, Birgi Tamersoy, Yifan Wu, Elena Balashova, Dorin Comaniciu
HIGHLIGHT: We present a novel framework that learns to predict human anatomy from body surface.

938, TITLE: Gibson Env: Real-World Perception for Embodied Agents
http://openaccess.thecvf.com/content_cvpr_2018/html/Xia_Gibson_Env_Real-World_CVPR_2018_paper.html
AUTHORS: Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, Silvio Savarese
HIGHLIGHT: In this paper, we investigate learning a real-world perception for active agents, propose Gibson virtual environment for this purpose, and showcase a set of learned complex locomotion abilities.

939, TITLE: Reinforcement Cutting-Agent Learning for Video Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Han_Reinforcement_Cutting-Agent_Learning_CVPR_2018_paper.html
AUTHORS: Junwei Han, Le Yang, Dingwen Zhang, Xiaojun Chang, Xiaodan Liang
HIGHLIGHT: In this paper, we formulate this problem as a Markov Decision Process, where agents are learned to segment object regions under a deep reinforcement learning framework.

940, TITLE: Feature Space Transfer for Data Augmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Liu_Feature_Space_Transfer_CVPR_2018_paper.html
AUTHORS: Bo Liu, Xudong Wang, Mandar Dixit, Roland Kwitt, Nuno Vasconcelos
HIGHLIGHT: Most notably, by using feature space transfer for data augmentation (w.r.t. pose and depth) on SUN-RGBD objects, we demonstrate considerable performance improvements on one/few-shot object recognition in a transfer learning setup, compared to current state-of-the-art methods.

941, TITLE: Analytic Expressions for Probabilistic Moments of PL-DNN With Gaussian Input
http://openaccess.thecvf.com/content_cvpr_2018/html/Bibi_Analytic_Expressions_for_CVPR_2018_paper.html
AUTHORS: Adel Bibi, Modar Alfadly, Bernard Ghanem
HIGHLIGHT: To this end, we derive in this paper exact analytic expressions for the first and second moments (mean and variance) of a small piecewise linear (PL) network (Affine, ReLU, Affine) subject to general Gaussian input.

942, TITLE: Detail-Preserving Pooling in Deep Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Saeedan_Detail-Preserving_Pooling_in_CVPR_2018_paper.html
AUTHORS: Faraz Saeedan, Nicolas Weber, Michael Goesele, Stefan Roth
HIGHLIGHT: In this paper, we aim to leverage recent results on image downscaling for the purposes of deep learning.

943, TITLE: Rethinking Feature Distribution for Loss Functions in Image Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Wan_Rethinking_Feature_Distribution_CVPR_2018_paper.html
AUTHORS: Weitao Wan, Yuanyi Zhong, Tianpeng Li, Jiansheng Chen
HIGHLIGHT: We propose a large-margin Gaussian Mixture (L-GM) loss for deep neural networks in classification tasks.

944, TITLE: Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions
http://openaccess.thecvf.com/content_cvpr_2018/html/Wu_Shift_A_Zero_CVPR_2018_paper.html
AUTHORS: Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, Kurt Keutzer
HIGHLIGHT: In this paper, we present a parameter-free, FLOP-free "shift" operation as an alternative to spatial convolutions.

- 945, TITLE: Sketch-a-Classifier: Sketch-Based Photo Classifier Generation
http://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Sketch-a-Classifier_Sketch-Based_Photo_CVPR_2018_paper.html
AUTHORS: Conghui Hu, Da Li, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales
HIGHLIGHT: In this paper we investigate an alternative approach of synthesizing image classifiers: almost directly from a user's imagination, via free-hand sketch.
- 946, TITLE: Light Field Intrinsic With a Deep Encoder-Decoder Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Alperovich_Light_Field_Intrinsic_CVPR_2018_paper.html
AUTHORS: Anna Alperovich, Ole Johannsen, Michael Strecke, Bastian Goldluecke
HIGHLIGHT: We present a fully convolutional autoencoder for light fields, which jointly encodes stacks of horizontal and vertical epipolar plane images through a deep network of residual layers.
- 947, TITLE: Learning Generative ConvNets via Multi-Grid Modeling and Sampling
http://openaccess.thecvf.com/content_cvpr_2018/html/Gao_Learning_Generative_ConvNets_CVPR_2018_paper.html
AUTHORS: Ruiqi Gao, Yang Lu, Junpei Zhou, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT: This paper proposes a multi-grid method for learning energy-based generative ConvNet models of images.
- 948, TITLE: Manifold Learning in Quotient Spaces
http://openaccess.thecvf.com/content_cvpr_2018/html/Mehr_Manifold_Learning_in_CVPR_2018_paper.html
AUTHORS: Éloi Mehr, André Lieutier, Fernando Sanchez Bermudez, Vincent Guitteny, Nicolas Thome, Matthieu Cord
HIGHLIGHT: In this paper we introduce a new autoencoder model for encoding and synthesis of 3D shapes.
- 949, TITLE: Learning Intelligent Dialogs for Bounding Box Annotation
http://openaccess.thecvf.com/content_cvpr_2018/html/Konyushkova_Learning_Intelligent_Dialogs_CVPR_2018_paper.html
AUTHORS: Ksenia Konyushkova, Jasper Uijlings, Christoph H. Lampert, Vittorio Ferrari
HIGHLIGHT: We introduce Intelligent Annotation Dialogs for bounding box annotation.
- 950, TITLE: Boosting Adversarial Attacks With Momentum
http://openaccess.thecvf.com/content_cvpr_2018/html/Dong_Boosting_Adversarial_Attacks_CVPR_2018_paper.html
AUTHORS: Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, Jianguo Li
HIGHLIGHT: To address this issue, we propose a broad class of momentum-based iterative algorithms to boost adversarial attacks.
- 951, TITLE: NISP: Pruning Networks Using Neuron Importance Score Propagation
http://openaccess.thecvf.com/content_cvpr_2018/html/Yu_NISP_Pruning_Networks_CVPR_2018_paper.html
AUTHORS: Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I. Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, Larry S. Davis
HIGHLIGHT: Based on our theoretical analysis, we propose the Neuron Importance Score Propagation (NISP) algorithm to propagate the importance scores of final responses to every neuron in the network.
- 952, TITLE: PointGrid: A Deep Network for 3D Shape Understanding
http://openaccess.thecvf.com/content_cvpr_2018/html/Le_PointGrid_A_Deep_CVPR_2018_paper.html
AUTHORS: Truc Le, Ye Duan
HIGHLIGHT: This paper presents a new deep learning architecture called PointGrid that is designed for 3D model recognition from unorganized point clouds.
- 953, TITLE: Tell Me Where to Look: Guided Attention Inference Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_Tell_Me_Where_CVPR_2018_paper.html
AUTHORS: Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, Yun Fu
HIGHLIGHT: In one common framework we address three shortcomings of previous approaches in modeling such attention maps: We (1) make attention maps an explicit and natural component of the end-to-end training for the first time, (2) provide self-guidance directly on these maps by exploring supervision from the network itself to improve them, and (3) seamlessly bridge the gap between using weak and extra supervision if available.
- 954, TITLE: 3D Semantic Segmentation With Submanifold Sparse Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Graham_3D_Semantic_Segmentation_CVPR_2018_paper.html
AUTHORS: Benjamin Graham, Martin Engelcke, Laurens van der Maaten
HIGHLIGHT: We introduce new sparse convolutional operations that are designed to process spatially-sparse data more efficiently, and use them to develop spatially-sparse convolutional networks.

- 955, TITLE: TOM-Net: Learning Transparent Object Matting From a Single Image
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_TOM-Net_Learning_Transparent_CVPR_2018_paper.html
AUTHORS: Guanying Chen, Kai Han, Kwan-Yee K. Wong
HIGHLIGHT: This paper addresses the problem of transparent object matting.
As no off-the-shelf dataset is available for transparent object matting, we create a large-scale synthetic dataset consisting of 178K images of transparent objects rendered in front of images sampled from the Microsoft COCO dataset.
We also collect a real dataset consisting of 876 samples using 14 transparent objects and 60 background images.
- 956, TITLE: Translating and Segmenting Multimodal Medical Volumes With Cycle- and Shape-Consistency Generative Adversarial Network
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Translating_and_Segmenting_CVPR_2018_paper.html
AUTHORS: Zizhao Zhang, Lin Yang, Yefeng Zheng
HIGHLIGHT: In this work, we propose a generic cross-modality synthesis approach with the following targets: 1) synthesizing realistic looking 3D images using unpaired training data, 2) ensuring consistent anatomical structures, which could be changed by geometric distortion in cross-modality synthesis and 3) improving volume segmentation by using synthetic data for modalities with limited training samples.
- 957, TITLE: An Unsupervised Learning Model for Deformable Medical Image Registration
http://openaccess.thecvf.com/content_cvpr_2018/html/Balakrishnan_An_Unsupervised_Learning_CVPR_2018_paper.html
AUTHORS: Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, Adrian V. Dalca
HIGHLIGHT: We present a fast learning-based algorithm for deformable, pairwise 3D medical image registration.
- 958, TITLE: Deep Lesion Graphs in the Wild: Relationship Learning and Organization of Significant Radiology Image Findings in a Diverse Large-Scale Lesion Database
http://openaccess.thecvf.com/content_cvpr_2018/html/Yan_Deep_Lesion_Graphs_CVPR_2018_paper.html
AUTHORS: Ke Yan, Xiaosong Wang, Le Lu, Ling Zhang, Adam P. Harrison, Mohammadhadi Bagheri, Ronald M. Summers
HIGHLIGHT: In this paper, we aim to organize and explore them by learning a deep feature representation for each lesion.
- 959, TITLE: Learning Distributions of Shape Trajectories From Longitudinal Datasets: A Hierarchical Model on a Manifold of Diffeomorphisms
http://openaccess.thecvf.com/content_cvpr_2018/html/Bone_Learning_Distributions_of_CVPR_2018_paper.html
AUTHORS: Alexandre Bône, Olivier Colliot, Stanley Durrleman
HIGHLIGHT: We propose a method to learn a distribution of shape trajectories from longitudinal data, i.e. the collection of individual objects repeatedly observed at multiple time-points.
- 960, TITLE: CNN Driven Sparse Multi-Level B-Spline Image Registration
http://openaccess.thecvf.com/content_cvpr_2018/html/Jiang_CNN_Driven_Sparse_CVPR_2018_paper.html
AUTHORS: Pingge Jiang, James A. Shackleford
HIGHLIGHT: To overcome these difficulties in determining B-spline grid configurations, this paper investigates the use of convolutional neural networks (CNNs) to learn and infer expressive sparse multi-grid configurations prior to B-spline coefficient optimization.
- 961, TITLE: Anatomical Priors in Convolutional Networks for Unsupervised Biomedical Segmentation
http://openaccess.thecvf.com/content_cvpr_2018/html/Dalca_Anatomical_Priors_in_CVPR_2018_paper.html
AUTHORS: Adrian V. Dalca, John Guttag, Mert R. Sabuncu
HIGHLIGHT: We introduce a generative probabilistic model that employs the learned prior through a convolutional neural network to compute segmentations in an unsupervised setting.
- 962, TITLE: 3D Registration of Curves and Surfaces Using Local Differential Information
http://openaccess.thecvf.com/content_cvpr_2018/html/Raposo_3D_Registration_of_CVPR_2018_paper.html
AUTHORS: Carolina Raposo, João P. Barreto
HIGHLIGHT: This article presents for the first time a global method for registering 3D curves with 3D surfaces without requiring an initialization.
- 963, TITLE: Weakly Supervised Learning of Single-Cell Feature Embeddings
http://openaccess.thecvf.com/content_cvpr_2018/html/Caicedo_Weakly_Supervised_Learning_CVPR_2018_paper.html
AUTHORS: Juan C. Caicedo, Claire McQuin, Allen Goodman, Shantanu Singh, Anne E. Carpenter
HIGHLIGHT: We therefore propose to train CNNs based on a weakly supervised approach, where the network aims to classify each treatment against all others.

964, TITLE: Guided Proofreading of Automatic Segmentations for Connectomics
http://openaccess.thecvf.com/content_cvpr_2018/html/Haehn_Guided_Proofreading_of_CVPR_2018_paper.html
AUTHORS: Daniel Haehn, Verena Kaynig, James Tompkin, Jeff W. Lichtman, Hanspeter Pfister
HIGHLIGHT: To aid error correction, we develop two classifiers that automatically recommend candidate merges and splits to the user.

965, TITLE: Wide Compression: Tensor Ring Nets
http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Wide_Compression_Tensor_CVPR_2018_paper.html
AUTHORS: Wenqi Wang, Yifan Sun, Brian Eriksson, Wenlin Wang, Vaneet Aggarwal
HIGHLIGHT: Inspired by the recent tensor ring factorization, we introduce Tensor Ring Networks (TR-Nets), which significantly compress both the fully connected layers and the convolutional layers of deep networks.

966, TITLE: Improvements to Context Based Self-Supervised Learning
http://openaccess.thecvf.com/content_cvpr_2018/html/Mundhenk_Improvements_to_Context_CVPR_2018_paper.html
AUTHORS: T. Nathan Mundhenk, Daniel Ho, Barry Y. Chen
HIGHLIGHT: We develop a set of methods to improve on the results of self-supervised learning using context.

967, TITLE: Learning Structure and Strength of CNN Filters for Small Sample Size Training
http://openaccess.thecvf.com/content_cvpr_2018/html/Keshari_Learning_Structure_and_CVPR_2018_paper.html
AUTHORS: Rohit Keshari, Mayank Vatsa, Richa Singh, Afzel Noore
HIGHLIGHT: To address this limitation, in this paper, we propose SSF-CNN which focuses on learning the "structure" and "strength" of filters.

968, TITLE: Boosting Self-Supervised Learning via Knowledge Transfer
http://openaccess.thecvf.com/content_cvpr_2018/html/Noroozi_Boosting_Self-Supervised_Learning_CVPR_2018_paper.html
AUTHORS: Mehdi Noroozi, Ananth Vinjimoor, Paolo Favaro, Hamed Pirsiavash
HIGHLIGHT: In this paper, we present a novel framework for self-supervised learning that overcomes limitations in designing and comparing different tasks, models, and data domains.

969, TITLE: The Power of Ensembles for Active Learning in Image Classification
http://openaccess.thecvf.com/content_cvpr_2018/html/Beluch_The_Power_of_CVPR_2018_paper.html
AUTHORS: William H. Beluch, Tim Genewein, Andreas Nürnberger, Jan M. Köhler
HIGHLIGHT: In this paper we investigate some recently proposed methods for active learning with high-dimensional data and convolutional neural network classifiers.

970, TITLE: Learning Compact Recurrent Neural Networks With Block-Term Tensor Decomposition
http://openaccess.thecvf.com/content_cvpr_2018/html/Ye_Learning_Compact_Recurrent_CVPR_2018_paper.html
AUTHORS: Jinmian Ye, Linnan Wang, Guangxi Li, Di Chen, Shandian Zhe, Xinqi Chu, Zenglin Xu
HIGHLIGHT: To overcome this problem, we propose a compact and flexible structure, namely Block-Term tensor decomposition, which greatly reduces the parameters of RNNs and improves their training efficiency.

971, TITLE: Spatially-Adaptive Filter Units for Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Tabernik_Spatially-Adaptive_Filter_Units_CVPR_2018_paper.html
AUTHORS: Domen Tabernik, Matej Kristan, Aleš Leonardis
HIGHLIGHT: In this paper we propose a novel displaced aggregation unit (DAU) that does not require hand-crafting.

972, TITLE: SO-Net: Self-Organizing Network for Point Cloud Analysis
http://openaccess.thecvf.com/content_cvpr_2018/html/Li_SO-Net_Self-Organizing_Network_CVPR_2018_paper.html
AUTHORS: Jiaxin Li, Ben M. Chen, Gim Hee Lee
HIGHLIGHT: This paper presents SO-Net, a permutation invariant architecture for deep learning with orderless point clouds.

973, TITLE: SGAN: An Alternative Training of Generative Adversarial Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Chavdarova_SGAN_An_Alternative_CVPR_2018_paper.html
AUTHORS: Tatjana Chavdarova, François Fleuret
HIGHLIGHT: This approach aims at increasing the chances that learning will not stop for the global pair, preventing both to be trapped in an unsatisfactory local minimum, or to face oscillations often observed in practice.

974, TITLE: SketchyGAN: Towards Diverse and Realistic Sketch to Image Synthesis
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_SketchyGAN_Towards_Diverse_CVPR_2018_paper.html

AUTHORS: Wengling Chen, James Hays
HIGHLIGHT: In this work, we propose a novel Generative Adversarial Network (GAN) approach that synthesizes plausible images from 50 categories including motorcycles, horses and couches.

975, TITLE: Explicit Loss-Error-Aware Quantization for Low-Bit Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhou_Explicit_Loss-Error-Aware_Quantization_CVPR_2018_paper.html
AUTHORS: Aojun Zhou, Anbang Yao, Kuan Wang, Yurong Chen
HIGHLIGHT: In this paper, we propose Explicit Loss-error-aware Quantization (ELQ), a new method that can train DNN models with very low-bit parameter values such as ternary and binary ones to approximate 32-bit floating-point counterparts without noticeable loss of predication accuracy.

976, TITLE: Towards Universal Representation for Unseen Action Recognition
http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_Towards_Universal_Representation_CVPR_2018_paper.html
AUTHORS: Yi Zhu, Yang Long, Yu Guan, Shawn Newsam, Ling Shao
HIGHLIGHT: Towards Universal Representation for Unseen Action Recognition

977, TITLE: Deep Image Prior
http://openaccess.thecvf.com/content_cvpr_2018/html/Ulyanov_Deep_Image_Prior_CVPR_2018_paper.html
AUTHORS: Dmitry Ulyanov, Andrea Vedaldi, Victor Lempitsky
HIGHLIGHT: In this paper, we show that, on the contrary, the structure of a generator network is sufficient to capture a great deal of low-level image statistics prior to any learning.

978, TITLE: ST-GAN: Spatial Transformer Generative Adversarial Networks for Image Compositing
http://openaccess.thecvf.com/content_cvpr_2018/html/Lin_ST-GAN_Spatial_Transformer_CVPR_2018_paper.html
AUTHORS: Chen-Hsuan Lin, Ersin Yumer, Oliver Wang, Eli Shechtman, Simon Lucey
HIGHLIGHT: To achieve this, we propose a novel Generative Adversarial Network (GAN) architecture that utilizes Spatial Transformer Networks (STNs) as the generator, which we call Spatial Transformer GANs (ST-GANs).

979, TITLE: CartoonGAN: Generative Adversarial Networks for Photo Cartoonization
http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_CartoonGAN_Generative_Adversarial_CVPR_2018_paper.html
AUTHORS: Yang Chen, Yu-Kun Lai, Yong-Jin Liu
HIGHLIGHT: In this paper, we propose a solution to transforming photos of real-world scenes into cartoon style images, which is valuable and challenging in computer vision and computer graphics.