

TITLE: Finding Task-Relevant Features for Few-Shot Learning by Category Traversal
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Finding_Task-Relevant_Features_for_Few-Shot_Learning_by_Category_Traversal_CVPR_2019_paper.html
AUTHORS: Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, Xiaogang Wang
HIGHLIGHT: In this work, we introduce a Category Traversal Module that can be inserted as a plug-and-play module into most metric-learning based few-shot learners.

TITLE: Edge-Labeling Graph Neural Network for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Edge-Labeling_Graph_Neural_Network_for_Few-Shot_Learning_CVPR_2019_paper.html
AUTHORS: Jongmin Kim, Taesup Kim, Sungwoong Kim, Chang D. Yoo
HIGHLIGHT: In this paper, we propose a novel edge-labeling graph neural network (EGNN), which adapts a deep neural network on the edge-labeling graph, for few-shot learning.

TITLE: Generating Classification Weights With GNN Denoising Autoencoders for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Gidaris_Generating_Classification_Weights_With_GNN_Denoising_Autoencoders_for_Few-Shot_Learning_CVPR_2019_paper.html
AUTHORS: Spyros Gidaris, Nikos Komodakis
HIGHLIGHT: Given an initial recognition model already trained on a set of base classes, the goal of this work is to develop a meta-model for few-shot learning.

TITLE: Kervolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Kervolutional_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Chen Wang, Jianfei Yang, Lihua Xie, Junsong Yuan
HIGHLIGHT: Existing works mainly leverage on the activation layers, which can only provide point-wise non-linearity.

TITLE: Why ReLU Networks Yield High-Confidence Predictions Far Away From the Training Data and How to Mitigate the Problem
http://openaccess.thecvf.com/content_CVPR_2019/html/Hein_Why_ReLU_Networks_Yield_High-Confidence_Predictions_Far_Away_From_the_CVPR_2019_paper.html
AUTHORS: Matthias Hein, Maksym Andriushchenko, Julian Bitterwolf
HIGHLIGHT: For bounded domains like images we propose a new robust optimization technique similar to adversarial training which enforces low confidence predictions far away from the training data.

TITLE: On the Structural Sensitivity of Deep Convolutional Networks to the Directions of Fourier Basis Functions
http://openaccess.thecvf.com/content_CVPR_2019/html/Tsuzuku_On_the_Structural_Sensitivity_of_Deep_Convolutional_Networks_to_the_CVPR_2019_paper.html
AUTHORS: Yusuke Tsuzuku, Issei Sato
HIGHLIGHT: As a byproduct of the analysis, we propose an algorithm to create shift-invariant universal adversarial perturbations available in black-box settings.

TITLE: Neural Rejuvenation: Improving Deep Network Training by Enhancing Computational Resource Utilization
http://openaccess.thecvf.com/content_CVPR_2019/html/Qiao_Neural_Rejuvenation_Improving_Deep_Network_Training_by_Enhancing_Computational_Resource_CVPR_2019_paper.html
AUTHORS: Siyuan Qiao, Zhe Lin, Jianming Zhang, Alan L. Yuille
HIGHLIGHT: In this paper, we study the problem of improving computational resource utilization of neural networks.

TITLE: Hardness-Aware Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Zheng_Hardness-Aware_Deep_Metric_Learning_CVPR_2019_paper.html
AUTHORS: Wenzhao Zheng, Zhaodong Chen, Jiwen Lu, Jie Zhou
HIGHLIGHT: This paper presents a hardness-aware deep metric learning (HDML) framework.

TITLE: Auto-DeepLab: Hierarchical Neural Architecture Search for Semantic Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Auto-DeepLab_Hierarchical_Neural_Architecture_Search_for_Semantic_Image_Segmentation_CVPR_2019_paper.html
AUTHORS: Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L. Yuille, Li Fei-Fei
HIGHLIGHT: In this paper, we study NAS for semantic image segmentation.

TITLE: Learning Loss for Active Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Yoo_Learning_Loss_for_Active_Learning_CVPR_2019_paper.html
AUTHORS: Donggeun Yoo, In So Kweon

HIGHLIGHT: In this paper, we propose a novel active learning method that is simple but task-agnostic, and works efficiently with the deep networks.

TITLE: Striking the Right Balance With Uncertainty

http://openaccess.thecvf.com/content_CVPR_2019/html/Khan_Striking_the_Right_Balance_With_Uncertainty_CVPR_2019_paper.html

AUTHORS: Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, Ling Shao

HIGHLIGHT: In this paper, we demonstrate that the Bayesian uncertainty estimates directly correlate with the rarity of classes and the difficulty level of individual samples.

TITLE: AutoAugment: Learning Augmentation Strategies From Data

http://openaccess.thecvf.com/content_CVPR_2019/html/Cubuk_AutoAugment_Learning_Augmentation_Strategies_From_Data_CVPR_2019_paper.html

AUTHORS: Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, Quoc V. Le

HIGHLIGHT: In this paper, we describe a simple procedure called AutoAugment to automatically search for improved data augmentation policies.

TITLE: Parsing R-CNN for Instance-Level Human Analysis

http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Parsing_R-CNN_for_Instance-Level_Human_Analysis_CVPR_2019_paper.html

AUTHORS: Lu Yang, Qing Song, Zhihui Wang, Ming Jiang

HIGHLIGHT: In this paper, we present an end-to-end pipeline for solving the instance-level human analysis, named Parsing R-CNN.

TITLE: Large Scale Incremental Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Large_Scale_Incremental_Learning_CVPR_2019_paper.html

AUTHORS: Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, Yun Fu

HIGHLIGHT: We propose a simple and effective method to address this data imbalance issue.

TITLE: TopNet: Structural Point Cloud Decoder

http://openaccess.thecvf.com/content_CVPR_2019/html/Tchapmi_TopNet_Structural_Point_Cloud_Decoder_CVPR_2019_paper.html

AUTHORS: Lyne P. Tchapmi, Vineet Kosaraju, Hamid Rezaatofghi, Ian Reid, Silvio Savarese

HIGHLIGHT: In this work, we propose a novel decoder that generates a structured point cloud without assuming any specific structure or topology on the underlying point set.

TITLE: Perceive Where to Focus: Learning Visibility-Aware Part-Level Features for Partial Person Re-Identification

http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Perceive_Where_to_Focus_Learning_Visibility-Aware_Part-Level_Features_for_Partial_CVPR_2019_paper.html

AUTHORS: Yifan Sun, Qin Xu, Yali Li, Chi Zhang, Yikang Li, Shengjin Wang, Jian Sun

HIGHLIGHT: We propose a Visibility-aware Part Model (VPM) for partial re-ID, which learns to perceive the visibility of regions through self-supervision.

TITLE: Meta-Transfer Learning for Few-Shot Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Meta-Transfer_Learning_for_Few-Shot_Learning_CVPR_2019_paper.html

AUTHORS: Qianru Sun, Yaoyao Liu, Tat-Seng Chua, Bernt Schiele

HIGHLIGHT: In this paper we propose a novel few-shot learning method called meta-transfer learning (MTL) which learns to adapt a deep NN for few shot learning tasks.

TITLE: Structured Binary Neural Networks for Accurate Image Classification and Semantic Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhuang_Structured_Binary_Neural_Networks_for_Accurate_Image_Classification_and_Semantic_CVPR_2019_paper.html

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

HIGHLIGHT: In this paper, we propose to train convolutional neural networks (CNNs) with both binarized weights and activations, leading to quantized models specifically for mobile devices with limited power capacity and computation resources.

TITLE: Deep RNN Framework for Visual Sequential Applications

http://openaccess.thecvf.com/content_CVPR_2019/html/Pang_Deep_RNN_Framework_for_Visual_Sequential_Applications_CVPR_2019_paper.html

AUTHORS: Bo Pang, Kaiwen Zha, Hanwen Cao, Chen Shi, Cewu Lu

HIGHLIGHT: To deal with this, we propose a new recurrent neural framework that can be stacked deep effectively.

TITLE: Graph-Based Global Reasoning Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Graph-Based_Global_Reasoning_Networks_CVPR_2019_paper.html
AUTHORS: Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, Yannis Kalantidis
HIGHLIGHT: In this work, we propose a new approach for reasoning globally in which a set of features are globally aggregated over the coordinate space and then projected to an interaction space where relational reasoning can be efficiently computed.

TITLE: SSN: Learning Sparse Switchable Normalization via SparsestMax
http://openaccess.thecvf.com/content_CVPR_2019/html/Shao_SSN_Learning_Sparse_Switchable_Normalization_via_SparsestMax_CVPR_2019_paper.html
AUTHORS: Wenqi Shao, Tianjian Meng, Jingyu Li, Ruimao Zhang, Yudian Li, Xiaogang Wang, Ping Luo
HIGHLIGHT: Normalization methods improve both optimization and generalization of ConvNets.

TITLE: Spherical Fractal Convolutional Neural Networks for Point Cloud Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Rao_Spherical_Fractal_Convolutional_Neural_Networks_for_Point_Cloud_Recognition_CVPR_2019_paper.html
AUTHORS: Yongming Rao, Jiwen Lu, Jie Zhou
HIGHLIGHT: We present a generic, flexible and 3D rotation invariant framework based on spherical symmetry for point cloud recognition.

TITLE: Learning to Generate Synthetic Data via Compositing
http://openaccess.thecvf.com/content_CVPR_2019/html/Tripathi_Learning_to_Generate_Synthetic_Data_via_Compositing_CVPR_2019_paper.html
AUTHORS: Shashank Tripathi, Siddhartha Chandra, Amit Agrawal, Ambrish Tyagi, James M. Rehg, Visesh Chari
HIGHLIGHT: We present a task-specific approach to synthetic data generation.

TITLE: Divide and Conquer the Embedding Space for Metric Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Sanakoyeu_Divide_and_Conquer_the_Embedding_Space_for_Metric_Learning_CVPR_2019_paper.html
AUTHORS: Artsiom Sanakoyeu, Vadim Tschernezki, Uta Buchler, Bjorn Ommer
HIGHLIGHT: In this work, we propose a novel easy-to-implement divide and conquer approach for deep metric learning, which significantly improves the state-of-the-art performance of metric learning.

TITLE: Latent Space Autoregression for Novelty Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Abati_Latent_Space_Autoregression_for_Novelty_Detection_CVPR_2019_paper.html
AUTHORS: Davide Abati, Angelo Porrello, Simone Calderara, Rita Cucchiara
HIGHLIGHT: In our proposal, we design a general unsupervised framework where we equip a deep autoencoder with a parametric density estimator that learns the probability distribution underlying the latent representations with an autoregressive procedure.

TITLE: Attending to Discriminative Certainty for Domain Adaptation
http://openaccess.thecvf.com/content_CVPR_2019/html/Kurmi_Attending_to_Discriminative_Certainty_for_Domain_Adaptation_CVPR_2019_paper.html
AUTHORS: Vinod Kumar Kurmi, Shanu Kumar, Vinay P. Namboodiri
HIGHLIGHT: In this paper, we aim to solve for unsupervised domain adaptation of classifiers where we have access to label information for the source domain while these are not available for a target domain.

TITLE: Feature Denoising for Improving Adversarial Robustness
http://openaccess.thecvf.com/content_CVPR_2019/html/Xie_Feature_Denoising_for_Improving_Adversarial_Robustness_CVPR_2019_paper.html
AUTHORS: Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L. Yuille, Kaiming He
HIGHLIGHT: This study suggests that adversarial perturbations on images lead to noise in the features constructed by these networks.

TITLE: Selective Kernel Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Selective_Kernel_Networks_CVPR_2019_paper.html
AUTHORS: Xiang Li, Wenhai Wang, Xiaolin Hu, Jian Yang

HIGHLIGHT: We propose a dynamic selection mechanism in CNNs that allows each neuron to adaptively adjust its receptive field size based on multiple scales of input information.

TITLE: On Implicit Filter Level Sparsity in Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Mehta_On_Implicit_Filter_Level_Sparsity_in_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Dushyant Mehta, Kwang In Kim, Christian Theobalt
HIGHLIGHT: We investigate filter level sparsity that emerges in convolutional neural networks (CNNs) which employ Batch Normalization and ReLU activation, and are trained with adaptive gradient descent techniques and L2 regularization or weight decay.

TITLE: FlowNet3D: Learning Scene Flow in 3D Point Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_FlowNet3D_Learning_Scene_Flow_in_3D_Point_Clouds_CVPR_2019_paper.html

AUTHORS: Xingyu Liu, Charles R. Qi, Leonidas J. Guibas
HIGHLIGHT: In this work, we propose a novel deep neural network named FlowNet3D that learns scene flow from point clouds in an end-to-end fashion.

TITLE: Scene Memory Transformer for Embodied Agents in Long-Horizon Tasks
http://openaccess.thecvf.com/content_CVPR_2019/html/Fang_Scene_Memory_Transformer_for_Embodied_Agents_in_Long-Horizon_Tasks_CVPR_2019_paper.html

AUTHORS: Kuan Fang, Alexander Toshev, Li Fei-Fei, Silvio Savarese
HIGHLIGHT: In this work, we propose a novel memory-based policy, named Scene Memory Transformer (SMT).

TITLE: Co-Occurrent Features in Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Co-Occurrent_Features_in_Semantic_Segmentation_CVPR_2019_paper.html

AUTHORS: Hang Zhang, Han Zhang, Chenguang Wang, Junyuan Xie
HIGHLIGHT: In this paper, we go beyond global context and explore the fine-grained representation using co-occurrent features by introducing Co-occurrent Feature Model, which predicts the distribution of co-occurrent features for a given target.

TITLE: Bag of Tricks for Image Classification with Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/He_Bag_of_Tricks_for_Image_Classification_with_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, Mu Li
HIGHLIGHT: In this paper, we will examine a collection of such refinements and empirically evaluate their impact on the final model accuracy through ablation study.

TITLE: Learning Channel-Wise Interactions for Binary Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Learning_Channel-Wise_Interactions_for_Binary_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Ziwei Wang, Jiwen Lu, Chenxin Tao, Jie Zhou, Qi Tian
HIGHLIGHT: In this paper, we propose a channel-wise interaction based binary convolutional neural network learning method (CI-BCNN) for efficient inference.

TITLE: Knowledge Adaptation for Efficient Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/He_Knowledge_Adaptation_for_Efficient_Semantic_Segmentation_CVPR_2019_paper.html

AUTHORS: Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, Youliang Yan
HIGHLIGHT: To tackle this dilemma, we propose a knowledge distillation method tailored for semantic segmentation to improve the performance of the compact FCNs with large overall stride.

TITLE: Parametric Noise Injection: Trainable Randomness to Improve Deep Neural Network Robustness Against Adversarial Attack
http://openaccess.thecvf.com/content_CVPR_2019/html/He_Parametric_Noise_Injection_Trainable_Randomness_to_Improve_Deep_Neural_Network_CVPR_2019_paper.html

AUTHORS: Zhezhi He, Adnan Siraj Rakin, Deliang Fan
HIGHLIGHT: In this work, we propose Parametric-Noise-Injection (PNI) which involves trainable Gaussian noise injection at each layer on either activation or weights through solving the Min-Max optimization problem, embedded with adversarial training.

TITLE: Gotta Adapt 'Em All: Joint Pixel and Feature-Level Domain Adaptation for Recognition in the Wild

http://openaccess.thecvf.com/content_CVPR_2019/html/Tran_Gotta_Adapt_Em_All_Joint_Pixel_and_Feature-Level_Domain_Adaptation_CVPR_2019_paper.html

AUTHORS: Luan Tran, Kihyuk Sohn, Xiang Yu, Xiaoming Liu, Manmohan Chandraker

HIGHLIGHT: At the feature level, inspired by insights from semi-supervised learning, we propose a classification-aware domain adversarial neural network that brings target examples into more classifiable regions of source domain.

TITLE: Understanding the Disharmony Between Dropout and Batch Normalization by Variance Shift

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Understanding_the_Disharmony_Between_Dropout_and_Batch_Normalization_by_Variance_CVPR_2019_paper.html

AUTHORS: Xiang Li, Shuo Chen, Xiaolin Hu, Jian Yang

HIGHLIGHT: Theoretically, we find that Dropout shifts the variance of a specific neural unit when we transfer the state of that network from training to test.

TITLE: Circulant Binary Convolutional Networks: Enhancing the Performance of 1-Bit DCNNs With Circulant Back Propagation

http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Circulant_Binary_Convolutional_Networks_Enhancing_the_Performance_of_1-Bit_DCNNs_CVPR_2019_paper.html

AUTHORS: Chunlei Liu, Wenrui Ding, Xin Xia, Baochang Zhang, Jiaxin Gu, Jianzhuang Liu, Rongrong Ji, David Doermann

HIGHLIGHT: To address this problem, we propose new circulant filters (CiFs) and a circulant binary convolution (CBCConv) to enhance the capacity of binarized convolutional features via our circulant back propagation (CBP).

TITLE: DeFusionNET: Defocus Blur Detection via Recurrently Fusing and Refining Multi-Scale Deep Features

http://openaccess.thecvf.com/content_CVPR_2019/html/Tang_DeFusionNET_Defocus_Blur_Detection_via_Recurrently_Fusing_and_Refining_Multi-Scale_CVPR_2019_paper.html

AUTHORS: Chang Tang, Xinzhong Zhu, Xinwang Liu, Lizhe Wang, Albert Zomaya

HIGHLIGHT: To deal with these issues, we propose a deep neural network which recurrently fuses and refines multi-scale deep features (DeFusionNet) for defocus blur detection.

TITLE: Deep Virtual Networks for Memory Efficient Inference of Multiple Tasks

http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Deep_Virtual_Networks_for_Memory_Efficient_Inference_of_Multiple_Tasks_CVPR_2019_paper.html

AUTHORS: Eunwoo Kim, Chanho Ahn, Philip H.S. Torr, Songhwai Oh

HIGHLIGHT: In particular, in this work we address the problem of memory efficient learning for multiple tasks.

TITLE: Universal Domain Adaptation

http://openaccess.thecvf.com/content_CVPR_2019/html/You_Universal_Domain_Adaptation_CVPR_2019_paper.html

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

HIGHLIGHT: To solve the universal domain adaptation problem, we propose Universal Adaptation Network (UAN).

TITLE: Improving Transferability of Adversarial Examples With Input Diversity

http://openaccess.thecvf.com/content_CVPR_2019/html/Xie_Improving_Transferability_of_Adversarial_Examples_With_Input_Diversity_CVPR_2019_paper.html

AUTHORS: Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, Alan L. Yuille

HIGHLIGHT: To this end, we propose to improve the transferability of adversarial examples by creating diverse input patterns.

TITLE: Sequence-To-Sequence Domain Adaptation Network for Robust Text Image Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Sequence-To-Sequence_Domain_Adaptation_Network_for_Robust_Text_Image_Recognition_CVPR_2019_paper.html

AUTHORS: Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, Dongxiang Zhang, Heng Tao Shen

HIGHLIGHT: In this paper, we develop a Sequence-to-Sequence Domain Adaptation Network (SSDAN) for robust text image recognition, which could exploit unsupervised sequence data by an attention-based sequence encoder-decoder network.

TITLE: Hybrid-Attention Based Decoupled Metric Learning for Zero-Shot Image Retrieval

http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Hybrid-Attention_Based_Decoupled_Metric_Learning_for_Zero-Shot_Image_Retrieval_CVPR_2019_paper.html

AUTHORS: Binghui Chen, Weihong Deng

HIGHLIGHT: We demonstrate the necessity of addressing the vital problems in ZSIR on the popular benchmarks, outperforming the state-of-the-art methods by a significant margin.

TITLE: Learning to Sample

http://openaccess.thecvf.com/content_CVPR_2019/html/Dovrat_Learning_to_Sample_CVPR_2019_paper.html
AUTHORS: Oren Dovrat, Itai Lang, Shai Avidan
HIGHLIGHT: To do that, we propose a deep network to simplify 3D point clouds.

TITLE: Few-Shot Learning via Saliency-Guided Hallucination of Samples
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Few-Shot_Learning_via_Saliency-Guided_Hallucination_of_Samples_CVPR_2019_paper.html
AUTHORS: Hongguang Zhang, Jing Zhang, Piotr Koniusz
HIGHLIGHT: In this paper, we follow the latter direction and present a novel data hallucination model.

TITLE: Variational Convolutional Neural Network Pruning
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Variational_Convolutional_Neural_Network_Pruning_CVPR_2019_paper.html
AUTHORS: Chenglong Zhao, Bingbing Ni, Jian Zhang, Qiwei Zhao, Wenjun Zhang, Qi Tian
HIGHLIGHT: We propose a variational Bayesian scheme for pruning convolutional neural networks in channel level.

TITLE: Towards Optimal Structured CNN Pruning via Generative Adversarial Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Lin_Towards_Optimal_Structured_CNN_Pruning_via_Generative_Adversarial_Learning_CVPR_2019_paper.html
AUTHORS: Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, David Doermann
HIGHLIGHT: In this paper, we propose an effective structured pruning approach that jointly prunes filters as well as other structures in an end-to-end manner.

TITLE: Exploiting Kernel Sparsity and Entropy for Interpretable CNN Compression
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Exploiting_Kernel_Sparsity_and_Entropy_for_Interpretable_CNN_Compression_CVPR_2019_paper.html
AUTHORS: Yuchao Li, Shaohui Lin, Baochang Zhang, Jianzhuang Liu, David Doermann, Yongjian Wu, Feiyue Huang, Rongrong Ji
HIGHLIGHT: In this paper, we investigate the problem of CNN compression from a novel interpretable perspective.

TITLE: Fully Quantized Network for Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Fully_Quantized_Network_for_Object_Detection_CVPR_2019_paper.html
AUTHORS: Rundong Li, Yan Wang, Feng Liang, Hongwei Qin, Junjie Yan, Rui Fan
HIGHLIGHT: In this paper, we demonstrate that many of these difficulties arise because of instability during the fine-tuning stage of the quantization process, and propose several novel techniques to overcome these instabilities.

TITLE: MnasNet: Platform-Aware Neural Architecture Search for Mobile
http://openaccess.thecvf.com/content_CVPR_2019/html/Tan_MnasNet_Platform-Aware_Neural_Architecture_Search_for_Mobile_CVPR_2019_paper.html
AUTHORS: Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, Quoc V. Le
HIGHLIGHT: In this paper, we propose an automated mobile neural architecture search (MNAS) approach, which explicitly incorporate model latency into the main objective so that the search can identify a model that achieves a good trade-off between accuracy and latency.

TITLE: Student Becoming the Master: Knowledge Amalgamation for Joint Scene Parsing, Depth Estimation, and More
http://openaccess.thecvf.com/content_CVPR_2019/html/Ye_Student_Becoming_the_Master_Knowledge_Amalgamation_for_Joint_Scene_Parsing_CVPR_2019_paper.html
AUTHORS: Jingwen Ye, Yixin Ji, Xinchao Wang, Kairi Ou, Dapeng Tao, Mingli Song
HIGHLIGHT: In this paper, we investigate a novel deep-model reusing task.

TITLE: Learning Video Representations From Correspondence Proposals
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Learning_Video_Representations_From_Correspondence_Proposals_CVPR_2019_paper.html
AUTHORS: Xingyu Liu, Joon-Young Lee, Hailin Jin
HIGHLIGHT: In this paper, we propose a novel neural network that learns video representations by aggregating information from potential correspondences.

TITLE: SiamRPN++: Evolution of Siamese Visual Tracking With Very Deep Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_SiamRPN_Evolution_of_Siamese_Visual_Tracking_With_Very_Deep_Networks_CVPR_2019_paper.html

AUTHORS: Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, Junjie Yan

HIGHLIGHT: In this work we prove the core reason comes from the lack of strict translation invariance.

TITLE: Sphere Generative Adversarial Network Based on Geometric Moment Matching

http://openaccess.thecvf.com/content_CVPR_2019/html/Park_Sphere_Generative_Adversarial_Network_Based_on_Geometric_Moment_Matching_CVPR_2019_paper.html

AUTHORS: Sung Woo Park, Junseok Kwon

HIGHLIGHT: In the paper, we mathematically prove the good properties of sphere GAN.

TITLE: Adversarial Attacks Beyond the Image Space

http://openaccess.thecvf.com/content_CVPR_2019/html/Zeng_Adversarial_Attacks_Beyond_the_Image_Space_CVPR_2019_paper.html

AUTHORS: Xiaohui Zeng, Chenxi Liu, Yu-Siang Wang, Weichao Qiu, Lingxi Xie, Yu-Wing Tai, Chi-Keung Tang, Alan L. Yuille

HIGHLIGHT: Though image-space adversaries can be interpreted as per-pixel albedo change, we verify that they cannot be well explained along these physically meaningful dimensions, which often have a non-local effect.

TITLE: Evading Defenses to Transferable Adversarial Examples by Translation-Invariant Attacks

http://openaccess.thecvf.com/content_CVPR_2019/html/Dong_Evading_Defenses_to_Transferable_Adversarial_Examples_by_Translation-Invariant_Attacks_CVPR_2019_paper.html

AUTHORS: Yinpeng Dong, Tianyu Pang, Hang Su, Jun Zhu

HIGHLIGHT: In this paper, we propose a translation-invariant attack method to generate more transferable adversarial examples against the defense models.

TITLE: Decoupling Direction and Norm for Efficient Gradient-Based L2 Adversarial Attacks and Defenses

http://openaccess.thecvf.com/content_CVPR_2019/html/Rony_Decoupling_Direction_and_Norm_for_Efficient_Gradient-Based_L2_Adversarial_Attacks_CVPR_2019_paper.html

AUTHORS: Jerome Rony, Luiz G. Hafemann, Luiz S. Oliveira, Ismail Ben Ayed, Robert Sabourin, Eric Granger

HIGHLIGHT: In this paper, an efficient approach is proposed to generate gradient-based attacks that induce misclassifications with low L2 norm, by decoupling the direction and the norm of the adversarial perturbation that is added to the image.

TITLE: A General and Adaptive Robust Loss Function

http://openaccess.thecvf.com/content_CVPR_2019/html/Barron_A_General_and_Adaptive_Robust_Loss_Function_CVPR_2019_paper.html

AUTHORS: Jonathan T. Barron

HIGHLIGHT: We present a generalization of the Cauchy/Lorentzian, Geman-McClure, Welsch/Leclerc, generalized Charbonnier, Charbonnier/pseudo-Huber/L1-L2, and L2 loss functions.

TITLE: Filter Pruning via Geometric Median for Deep Convolutional Neural Networks Acceleration

http://openaccess.thecvf.com/content_CVPR_2019/html/He_Filter_Pruning_via_Geometric_Median_for_Deep_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, Yi Yang

HIGHLIGHT: To solve this problem, we propose a novel filter pruning method, namely Filter Pruning via Geometric Median (FPGM), to compress the model regardless of those two requirements.

TITLE: Learning to Quantize Deep Networks by Optimizing Quantization Intervals With Task Loss

http://openaccess.thecvf.com/content_CVPR_2019/html/Jung_Learning_to_Quantize_Deep_Networks_by_Optimizing_Quantization_Intervals_With_CVPR_2019_paper.html

AUTHORS: Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, Changkyu Choi

HIGHLIGHT: To tackle this problem, we propose to learn to quantize activations and weights via a trainable quantizer that transforms and discretizes them.

TITLE: Not All Areas Are Equal: Transfer Learning for Semantic Segmentation via Hierarchical Region Selection

http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Not_All_Areas_Are_Equal_Transfer_Learning_for_Semantic_Segmentation_CVPR_2019_paper.html

AUTHORS: Ruoqi Sun, Xinge Zhu, Chongruo Wu, Chen Huang, Jianping Shi, Lizhuang Ma

HIGHLIGHT: In this paper, we consider transfer learning for semantic segmentation that aims to mitigate the gap between abundant synthetic data (source domain) and limited real data (target domain).

TITLE: Unsupervised Learning of Dense Shape Correspondence
http://openaccess.thecvf.com/content_CVPR_2019/html/Halimi_Unsupervised_Learning_of_Dense_Shape_Correspondence_CVPR_2019_paper.html
AUTHORS: Oshri Halimi, Or Litany, Emanuele Rodola, Alex M. Bronstein, Ron Kimmel
HIGHLIGHT: We introduce the first completely unsupervised correspondence learning approach for deformable 3D shapes.

TITLE: Unsupervised Visual Domain Adaptation: A Deep Max-Margin Gaussian Process Approach
http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Unsupervised_Visual_Domain_Adaptation_A_Deep_Max-Margin_Gaussian_Process_Approach_CVPR_2019_paper.html
AUTHORS: Minyoung Kim, Pritish Sahu, Behnam Gholami, Vladimir Pavlovic
HIGHLIGHT: We formulate a learning objective that effectively influences the posterior to minimize the maximum discrepancy.

TITLE: Balanced Self-Paced Learning for Generative Adversarial Clustering Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Ghasedi_Balanced_Self-Paced_Learning_for_Generative_Adversarial_Clustering_Network_CVPR_2019_paper.html
AUTHORS: Kamran Ghasedi, Xiaoqian Wang, Cheng Deng, Heng Huang
HIGHLIGHT: In this paper, we propose a deep Generative Adversarial Clustering Network (ClusterGAN), which tackles the problems of training of deep clustering models in unsupervised manner.

TITLE: A Style-Based Generator Architecture for Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Karras_A_Style-Based_Generator_Architecture_for_Generative_Adversarial_Networks_CVPR_2019_paper.html
AUTHORS: Tero Karras, Samuli Laine, Timo Aila
HIGHLIGHT: To quantify interpolation quality and disentanglement, we propose two new, automated methods that are applicable to any generator architecture.

TITLE: Parallel Optimal Transport GAN
http://openaccess.thecvf.com/content_CVPR_2019/html/Avraham_Parallel_Optimal_Transport_GAN_CVPR_2019_paper.html
AUTHORS: Gil Avraham, Yan Zuo, Tom Drummond
HIGHLIGHT: To address these issues, we introduce an additional regularisation term which performs optimal transport in parallel within a low dimensional representation space.

TITLE: Reversible GANs for Memory-Efficient Image-To-Image Translation
http://openaccess.thecvf.com/content_CVPR_2019/html/van_der_Ouderaa_Reversible_GANs_for_Memory-Efficient_Image-To-Image_Translation_CVPR_2019_paper.html
AUTHORS: Tycho F.A. van der Ouderaa, Daniel E. Worrall
HIGHLIGHT: We extend this framework by exploring approximately invertible architectures which are well suited to these losses.

TITLE: Sensitive-Sample Fingerprinting of Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/He_Sensitive-Sample_Fingerprinting_of_Deep_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Zecheng He, Tianwei Zhang, Ruby Lee
HIGHLIGHT: In this paper, we propose a novel and practical methodology to verify the integrity of remote deep learning models, with only black-box access to the target models.

TITLE: Soft Labels for Ordinal Regression
http://openaccess.thecvf.com/content_CVPR_2019/html/Diaz_Soft_Labels_for_Ordinal_Regression_CVPR_2019_paper.html
AUTHORS: Raul Diaz, Amit Marathe
HIGHLIGHT: We present a simple and effective method that constrains these relationships among categories by seamlessly incorporating metric penalties into ground truth label representations.

TITLE: Local to Global Learning: Gradually Adding Classes for Training Deep Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Cheng_Local_to_Global_Learning_Gradually_Adding_Classes_for_Training_Deep_CVPR_2019_paper.html
AUTHORS: Hao Cheng, Dongze Lian, Bowen Deng, Shenghua Gao, Tao Tan, Yanlin Geng
HIGHLIGHT: In this paper, we incorporate the idea of LGL into the learning objective of DNNs and explain why LGL works better from an information-theoretic perspective.

TITLE: What Does It Mean to Learn in Deep Networks? And, How Does One Detect Adversarial Attacks?
http://openaccess.thecvf.com/content_CVPR_2019/html/Comeanu_What_Does_It_Mean_to_Learn_in_Deep_Networks_And_CVPR_2019_paper.html
AUTHORS: Ciprian A. Comeanu, Meysam Madadi, Sergio Escalera, Aleix M. Martinez
HIGHLIGHT: Here, we derive a novel approach to define what it means to learn in deep networks, and how to use this knowledge to detect adversarial attacks.

TITLE: Handwriting Recognition in Low-Resource Scripts Using Adversarial Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Bhunia_Handwriting_Recognition_in_Low-Resource_Scripts_Using_Adversarial_Learning_CVPR_2019_paper.html
AUTHORS: Ayan Kumar Bhunia, Abhirup Das, Ankan Kumar Bhunia, Perla Sai Raj Kishore, Partha Pratim Roy
HIGHLIGHT: We propose an Adversarial Feature Deformation Module (AFDM) that learns ways to elastically warp extracted features in a scalable manner.

TITLE: Adversarial Defense Through Network Profiling Based Path Extraction
http://openaccess.thecvf.com/content_CVPR_2019/html/Qiu_Adversarial_Defense_Through_Network_Profiling_Based_Path_Extraction_CVPR_2019_paper.html
AUTHORS: Yuxian Qiu, Jingwen Leng, Cong Guo, Quan Chen, Chao Li, Minyi Guo, Yuhao Zhu
HIGHLIGHT: This work proposes a profiling-based method to decompose the DNN models to different functional blocks, which lead to the effective path as a new approach to exploring DNNs' internal organization.

TITLE: RENAS: Reinforced Evolutionary Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_RENAS_Reinforced_Evolutionary_Neural_Architecture_Search_CVPR_2019_paper.html
AUTHORS: Yukang Chen, Gaofeng Meng, Qian Zhang, Shiming Xiang, Chang Huang, Lisen Mu, Xinggang Wang
HIGHLIGHT: To address this issue, we propose the Reinforced Evolutionary Neural Architecture Search (RENAS), which is an evolutionary method with reinforced mutation for NAS.

TITLE: Co-Occurrence Neural Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Shevlev_Co-Occurrence_Neural_Network_CVPR_2019_paper.html
AUTHORS: Irina Shevlev, Shai Avidan
HIGHLIGHT: We show how to train the filter as part of the network and report results on several data sets.

TITLE: SpotTune: Transfer Learning Through Adaptive Fine-Tuning
http://openaccess.thecvf.com/content_CVPR_2019/html/Guo_SpotTune_Transfer_Learning_Through_Adaptive_Fine-Tuning_CVPR_2019_paper.html
AUTHORS: Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, Rogerio Feris
HIGHLIGHT: In this paper, we propose an adaptive fine-tuning approach, called SpotTune, which finds the optimal fine-tuning strategy per instance for the target data.

TITLE: Signal-To-Noise Ratio: A Robust Distance Metric for Deep Metric Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Yuan_Signal-To-Noise_Ratio_A_Robust_Distance_Metric_for_Deep_Metric_Learning_CVPR_2019_paper.html
AUTHORS: Tongtong Yuan, Weihong Deng, Jian Tang, Yanan Tang, Binghui Chen
HIGHLIGHT: In this paper, different from the approaches on learning the loss structures, we propose a robust SNR distance metric based on Signal-to-Noise Ratio (SNR) for measuring the similarity of image pairs for deep metric learning.

TITLE: Detection Based Defense Against Adversarial Examples From the Steganalysis Point of View
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Detection_Based_Defense_Against_Adversarial_Examples_From_the_Steganalysis_Point_CVPR_2019_paper.html
AUTHORS: Jiayang Liu, Weiming Zhang, Yiwei Zhang, Dongdong Hou, Yujia Liu, Hongyue Zha, Nenghai Yu
HIGHLIGHT: In this paper, we point out that steganalysis can be applied to adversarial examples detection, and propose a method to enhance steganalysis features by estimating the probability of modifications caused by adversarial attacks.

TITLE: HetConv: Heterogeneous Kernel-Based Convolutions for Deep CNNs
http://openaccess.thecvf.com/content_CVPR_2019/html/Singh_HetConv_Heterogeneous_Kernel-Based_Convolutions_for_Deep_CNNs_CVPR_2019_paper.html
AUTHORS: Pravendra Singh, Vinay Kumar Verma, Piyush Rai, Vinay P. Namboodiri
HIGHLIGHT: We present a novel deep learning architecture in which the convolution operation leverages heterogeneous kernels.

TITLE: Strike (With) a Pose: Neural Networks Are Easily Fooled by Strange Poses of Familiar Objects
http://openaccess.thecvf.com/content_CVPR_2019/html/Alcorn_Strike_With_a_Pose_Neural_Networks_Are_Easily_Fooled_by_CVPR_2019_paper.html

AUTHORS: Michael A. Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, Anh Nguyen

HIGHLIGHT: In this paper, we present a framework for discovering DNN failures that harnesses 3D renderers and 3D models.

TITLE: Blind Geometric Distortion Correction on Images Through Deep Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Blind_Geometric_Distortion_Correction_on_Images_Through_Deep_Learning_CVPR_2019_paper.html

AUTHORS: Xiaoyu Li, Bo Zhang, Pedro V. Sander, Jing Liao

HIGHLIGHT: We propose the first general framework to automatically correct different types of geometric distortion in a single input image.

TITLE: Instance-Level Meta Normalization

http://openaccess.thecvf.com/content_CVPR_2019/html/Jia_Instance-Level_Meta_Normalization_CVPR_2019_paper.html

AUTHORS: Songhao Jia, Ding-Jie Chen, Hwann-Tzong Chen

HIGHLIGHT: This paper presents a normalization mechanism called Instance-Level Meta Normalization (ILM Norm) to address a learning-to-normalize problem.

TITLE: Iterative Normalization: Beyond Standardization Towards Efficient Whitening

http://openaccess.thecvf.com/content_CVPR_2019/html/Huang_Iterative_Normalization_Beyond_Standardization_Towards_Efficient_Whitening_CVPR_2019_paper.html

AUTHORS: Lei Huang, Yi Zhou, Fan Zhu, Li Liu, Ling Shao

HIGHLIGHT: We propose Iterative Normalization (IterNorm), which employs Newton's iterations for much more efficient whitening, while simultaneously avoiding the eigen-decomposition.

TITLE: On Learning Density Aware Embeddings

http://openaccess.thecvf.com/content_CVPR_2019/html/Ghosh_On_Learning_Density_Aware_Embeddings_CVPR_2019_paper.html

AUTHORS: Soumyadeep Ghosh, Richa Singh, Mayank Vatsa

HIGHLIGHT: In this paper, a novel noise tolerant deep metric learning algorithm is proposed.

TITLE: Contrastive Adaptation Network for Unsupervised Domain Adaptation

http://openaccess.thecvf.com/content_CVPR_2019/html/Kang_Contrastive_Adaptation_Network_for_Unsupervised_Domain_Adaptation_CVPR_2019_paper.html

AUTHORS: Guoliang Kang, Lu Jiang, Yi Yang, Alexander G. Hauptmann

HIGHLIGHT: To address this issue, this paper proposes Contrastive Adaptation Network (CAN) optimizing a new metric which explicitly models the intra-class domain discrepancy and the inter-class domain discrepancy.

TITLE: LP-3DCNN: Unveiling Local Phase in 3D Convolutional Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Kumawat_LP-3DCNN_Unveiling_Local_Phase_in_3D_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Sudhakar Kumawat, Shanmuganathan Raman

HIGHLIGHT: To address these issues, we propose Rectified Local Phase Volume (ReLPV) block, an efficient alternative to the standard 3D convolutional layer.

TITLE: Attribute-Driven Feature Disentangling and Temporal Aggregation for Video Person Re-Identification

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Attribute-Driven_Feature_Disentangling_and_Temporal_Aggregation_for_Video_Person_Re-Identification_CVPR_2019_paper.html

AUTHORS: Yiru Zhao, Xu Shen, Zhongming Jin, Hongtao Lu, Xian-sheng Hua

HIGHLIGHT: In this paper, we propose an attribute-driven method for feature disentangling and frame re-weighting.

TITLE: Binary Ensemble Neural Network: More Bits per Network or More Networks per Bit?

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhu_Binary_Ensemble_Neural_Network_More_Bits_per_Network_or_More_CVPR_2019_paper.html

AUTHORS: Shilin Zhu, Xin Dong, Hao Su

HIGHLIGHT: Inspired by this investigation, we propose the Binary Ensemble Neural Network (BENN) which leverages ensemble methods to improve the performance of BNNs with limited efficiency cost.

TITLE: Distilling Object Detectors With Fine-Grained Feature Imitation

http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Distilling_Object_Detectors_With_Fine-Grained_Feature_Imitation_CVPR_2019_paper.html

AUTHORS: Tao Wang, Li Yuan, Xiaopeng Zhang, Jiashi Feng
HIGHLIGHT: To address the challenge of distilling knowledge in detection model, we propose a fine-grained feature imitation method exploiting the cross-location discrepancy of feature response.

TITLE: Centripetal SGD for Pruning Very Deep Convolutional Networks With Complicated Structure
http://openaccess.thecvf.com/content_CVPR_2019/html/Ding_Centripetal_SGD_for_Pruning_Very_Deep_Convolutional_Networks_With_Complicated_CVPR_2019_paper.html

AUTHORS: Xiaohan Ding, Guiguang Ding, Yuchen Guo, Jungong Han
HIGHLIGHT: To this end, we propose Centripetal SGD (C-SGD), a novel optimization method, which can train several filters to collapse into a single point in the parameter hyperspace.

TITLE: Knockoff Nets: Stealing Functionality of Black-Box Models
http://openaccess.thecvf.com/content_CVPR_2019/html/Orekondy_Knockoff_Nets_Stealing_Functionality_of_Black-Box_Models_CVPR_2019_paper.html

AUTHORS: Tribhuvanesh Orekondy, Bernt Schiele, Mario Fritz
HIGHLIGHT: In this work, we ask to what extent can an adversary steal functionality of such "victim" models based solely on blackbox interactions: image in, predictions out.

TITLE: Bidirectional Learning for Domain Adaptation of Semantic Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Bidirectional_Learning_for_Domain_Adaptation_of_Semantic_Segmentation_CVPR_2019_paper.html

AUTHORS: Yunsheng Li, Lu Yuan, Nuno Vasconcelos
HIGHLIGHT: In this paper, we propose a novel bidirectional learning framework for domain adaptation of segmentation.

TITLE: Enhanced Bayesian Compression via Deep Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Yuan_Enhanced_Bayesian_Compression_via_Deep_Reinforcement_Learning_CVPR_2019_paper.html

AUTHORS: Xin Yuan, Liangliang Ren, Jiwen Lu, Jie Zhou
HIGHLIGHT: In this paper, we propose an Enhanced Bayesian Compression method to flexibly compress the deep networks via reinforcement learning.

TITLE: Strong-Weak Distribution Alignment for Adaptive Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Saito_Strong-Weak_Distribution_Alignment_for_Adaptive_Object_Detection_CVPR_2019_paper.html

AUTHORS: Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, Kate Saenko
HIGHLIGHT: We propose an approach for unsupervised adaptation of object detectors from label-rich to label-poor domains which can significantly reduce annotation costs associated with detection.

TITLE: MFAS: Multimodal Fusion Architecture Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Perez-Rua_MFAS_Multimodal_Fusion_Architecture_Search_CVPR_2019_paper.html

AUTHORS: Juan-Manuel Perez-Rua, Valentin Vielzeuf, Stephane Pateux, Moez Baccouche, Frederic Jurie
HIGHLIGHT: We propose a novel and generic search space that spans a large number of possible fusion architectures.

TITLE: Disentangling Adversarial Robustness and Generalization
http://openaccess.thecvf.com/content_CVPR_2019/html/Stutz_Disentangling_Adversarial_Robustness_and_Generalization_CVPR_2019_paper.html

AUTHORS: David Stutz, Matthias Hein, Bernt Schiele
HIGHLIGHT: To confirm our claims, we present extensive experiments on synthetic data (with known manifold) as well as on EMNIST, Fashion-MNIST and CelebA.

TITLE: ShieldNets: Defending Against Adversarial Attacks Using Probabilistic Adversarial Robustness
http://openaccess.thecvf.com/content_CVPR_2019/html/Theagarajan_ShieldNets_Defending_Against_Adversarial_Attacks_Using_Probabilistic_Adversarial_Robustness_CVPR_2019_paper.html

AUTHORS: Rajkumar Theagarajan, Ming Chen, Bir Bhanu, Jing Zhang
HIGHLIGHT: ShieldNet is implemented as a demonstration of PAR in this work by using PixelCNN.

TITLE: Deeply-Supervised Knowledge Synergy
http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Deeply-Supervised_Knowledge_Synergy_CVPR_2019_paper.html

AUTHORS: Dawei Sun, Anbang Yao, Aojun Zhou, Hao Zhao

HIGHLIGHT: In this paper, we propose Deeply-supervised Knowledge Synergy (DKS), a new method aiming to train CNNs with improved generalization ability for image classification tasks without introducing extra computational cost during inference.

TITLE: Dual Residual Networks Leveraging the Potential of Paired Operations for Image Restoration
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Dual_Residual_Networks_Leveraging_the_Potential_of_Paired_Operations_for_CVPR_2019_paper.html

AUTHORS: Xing Liu, Masanori Suganuma, Zhun Sun, Takayuki Okatani

HIGHLIGHT: In this paper, we study design of deep neural networks for tasks of image restoration.

TITLE: Probabilistic End-To-End Noise Correction for Learning With Noisy Labels

http://openaccess.thecvf.com/content_CVPR_2019/html/Yi_Probabilistic_End-To-End_Noise_Correction_for_Learning_With_Noisy_Labels_CVPR_2019_paper.html

AUTHORS: Kun Yi, Jianxin Wu

HIGHLIGHT: To address this problem, we propose an end-to-end framework called PENCIL, which can update both network parameters and label estimations as label distributions.

TITLE: Attention-Guided Unified Network for Panoptic Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Attention-Guided_Unified_Network_for_Panoptic_Segmentation_CVPR_2019_paper.html

AUTHORS: Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, Xingang Wang

HIGHLIGHT: Existing methods mostly dealt with these two problems separately, but in this paper, we reveal the underlying relationship between them, in particular, FG objects provide complementary cues to assist BG understanding.

TITLE: NAS-FPN: Learning Scalable Feature Pyramid Architecture for Object Detection

http://openaccess.thecvf.com/content_CVPR_2019/html/Ghiasi_NAS-FPN_Learning_Scalable_Feature_Pyramid_Architecture_for_Object_Detection_CVPR_2019_paper.html

AUTHORS: Golnaz Ghiasi, Tsung-Yi Lin, Quoc V. Le

HIGHLIGHT: Here we aim to learn a better architecture of feature pyramid network for object detection.

TITLE: OICSR: Out-In-Channel Sparsity Regularization for Compact Deep Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_OICSR_Out-In-Channel_Sparsity_Regularization_for_Compact_Deep_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Jiashi Li, Qi Qi, Jingyu Wang, Ce Ge, Yujian Li, Zhangzhang Yue, Haifeng Sun

HIGHLIGHT: Our method is comprehensively evaluated with various CNN architectures including CifarNet, AlexNet, ResNet, DenseNet and PreActSeNet on CIFAR-10, CIFAR-100 and ImageNet-1K datasets.

TITLE: Semantically Aligned Bias Reducing Zero Shot Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Paul_Semantically_Aligned_Bias_Reducing_Zero_Shot_Learning_CVPR_2019_paper.html

AUTHORS: Akanksha Paul, Narayanan C. Krishnan, Prateek Munjal

HIGHLIGHT: In this work, we propose a novel approach, Semantically Aligned Bias Reducing (SABR) ZSL, which focuses on solving both the problems.

TITLE: Feature Space Perturbations Yield More Transferable Adversarial Examples

http://openaccess.thecvf.com/content_CVPR_2019/html/Inkawhich_Feature_Space_Perturbations_Yield_More_Transferable_Adversarial_Examples_CVPR_2019_paper.html

AUTHORS: Nathan Inkawhich, Wei Wen, Hai (Helen) Li, Yiran Chen

HIGHLIGHT: This work describes a transfer-based blackbox targeted adversarial attack of deep feature space representations that also provides insights into cross-model class representations of deep CNNs.

TITLE: IGE-Net: Inverse Graphics Energy Networks for Human Pose Estimation and Single-View Reconstruction

http://openaccess.thecvf.com/content_CVPR_2019/html/Jack_IGE-Net_Inverse_Graphics_Energy_Networks_for_Human_Pose_Estimation_and_CVPR_2019_paper.html

AUTHORS: Dominic Jack, Frederic Maire, Sareh Shirazi, Anders Eriksson

HIGHLIGHT: We propose using a deep-learning based energy minimization framework to learn a consistency measure between 2D observations and a proposed world model, and demonstrate that this framework can be trained end-to-end to produce consistent and realistic inferences.

TITLE: Accelerating Convolutional Neural Networks via Activation Map Compression

http://openaccess.thecvf.com/content_CVPR_2019/html/Georgiadis_Accelerating_Convolutional_Neural_Networks_via_Activation_Map_Compression_CVPR_2019_paper.html

AUTHORS: Georgios Georgiadis
HIGHLIGHT: Towards this end, we propose a three-stage compression and acceleration pipeline that sparsifies, quantizes and entropy encodes activation maps of Convolutional Neural Networks.

TITLE: Knowledge Distillation via Instance Relationship Graph
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Knowledge_Distillation_via_Instance_Relationship_Graph_CVPR_2019_paper.html

AUTHORS: Yufan Liu, Jiajiong Cao, Bing Li, Chunfeng Yuan, Weiming Hu, Yangxi Li, Yunqiang Duan
HIGHLIGHT: In this paper, a novel Instance Relationship Graph (IRG) is proposed for knowledge distillation.

TITLE: PPGNet: Learning Point-Pair Graph for Line Segment Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_PPGNet_Learning_Point-Pair_Graph_for_Line_Segment_Detection_CVPR_2019_paper.html

AUTHORS: Ziheng Zhang, Zhengxin Li, Ning Bi, Jia Zheng, Jinlei Wang, Kun Huang, Weixin Luo, Yanyu Xu, Shenghua Gao
HIGHLIGHT: In this paper, we present a novel framework to detect line segments in man-made environments.

TITLE: Building Detail-Sensitive Semantic Segmentation Networks With Polynomial Pooling
http://openaccess.thecvf.com/content_CVPR_2019/html/Wei_Building_Detail-Sensitive_Semantic_Segmentation_Networks_With_Polynomial_Pooling_CVPR_2019_paper.html

AUTHORS: Zhen Wei, Jingyi Zhang, Li Liu, Fan Zhu, Fumin Shen, Yi Zhou, Si Liu, Yao Sun, Ling Shao
HIGHLIGHT: In this work, we propose a polynomial pooling (P-pooling) function that finds an intermediate form between max and average pooling to provide an optimally balanced and self-adjusted pooling strategy for semantic segmentation.

TITLE: Variational Bayesian Dropout With a Hierarchical Prior
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Variational_Bayesian_Dropout_With_a_Hierarchical_Prior_CVPR_2019_paper.html

AUTHORS: Yuhang Liu, Wenyong Dong, Lei Zhang, Dong Gong, Qinfeng Shi
HIGHLIGHT: To address this problem, we present a new generalization of Gaussian dropout, termed variational Bayesian dropout (VBD), which turns to exploit a hierarchical prior on the network weights and infer a new joint posterior.

TITLE: AANet: Attribute Attention Network for Person Re-Identifications
http://openaccess.thecvf.com/content_CVPR_2019/html/Tay_AANet_Attribute_Attention_Network_for_Person_Re-Identifications_CVPR_2019_paper.html

AUTHORS: Chiat-Pin Tay, Sharmili Roy, Kim-Hui Yap
HIGHLIGHT: This paper proposes Attribute Attention Network (AANet), a new architecture that integrates person attributes and attribute attention maps into a classification framework to solve the person re-identification (re-ID) problem.

TITLE: Overcoming Limitations of Mixture Density Networks: A Sampling and Fitting Framework for Multimodal Future Prediction
http://openaccess.thecvf.com/content_CVPR_2019/html/Makansi_Overcoming_Limitations_of_Mixture_Density_Networks_A_Sampling_and_Fitting_CVPR_2019_paper.html

AUTHORS: Osama Makansi, Eddy Ilg, Ozgun Cicek, Thomas Brox
HIGHLIGHT: In this work, we present an approach that involves the prediction of several samples of the future with a winner-takes-all loss and iterative grouping of samples to multiple modes.

TITLE: A Main/Subsidiary Network Framework for Simplifying Binary Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_A_MainSubsidiary_Network_Framework_for_Simplifying_Binary_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Yinghao Xu, Xin Dong, Yudian Li, Hao Su
HIGHLIGHT: In this paper, we, for the first time, define the filter-level pruning problem for binary neural networks, which cannot be solved by simply migrating existing structural pruning methods for full-precision models.

TITLE: PointNetLK: Robust & Efficient Point Cloud Registration Using PointNet
http://openaccess.thecvf.com/content_CVPR_2019/html/Aoki_PointNetLK_Robust_Efficient_Point_Cloud_Registration_Using_PointNet_CVPR_2019_paper.html

AUTHORS: Yasuhiro Aoki, Hunter Goforth, Rangaprasad Arun Srivatsan, Simon Lucey
HIGHLIGHT: In this paper we argue that PointNet itself can be thought of as a learnable "imaging" function.

TITLE: A Flexible Convolutional Solver for Fast Style Transfers

http://openaccess.thecvf.com/content_CVPR_2019/html/Puy_A_Flexible_Convolutional_Solver_for_Fast_Style_Transfers_CVPR_2019_paper.html

AUTHORS: Gilles Puy, Patrick Perez

HIGHLIGHT: We propose a new flexible deep convolutional neural network (convnet) to perform fast neural style transfers.

TITLE: Cross Domain Model Compression by Structurally Weight Sharing

http://openaccess.thecvf.com/content_CVPR_2019/html/Gao_Cross_Domain_Model_Compression_by_Structurally_Weight_Sharing_CVPR_2019_paper.html

AUTHORS: Shangqian Gao, Cheng Deng, Heng Huang

HIGHLIGHT: In this paper, thus, we propose a new robust cross domain model compression method.

TITLE: TraVeLGAN: Image-To-Image Translation by Transformation Vector Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Amodio_TraVeLGAN_Image-To-Image_Translation_by_Transformation_Vector_Learning_CVPR_2019_paper.html

AUTHORS: Matthew Amodio, Smita Krishnaswamy

HIGHLIGHT: For this purpose, we introduce a novel GAN based on preserving intra-domain vector transformations in a latent space learned by a siamese network.

TITLE: Deep Robust Subjective Visual Property Prediction in Crowdsourcing

http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_Deep_Robust_Subjective_Visual_Property_Prediction_in_Crowdsourcing_CVPR_2019_paper.html

AUTHORS: Qianqian Xu, Zhiyong Yang, Yangbangyan Jiang, Xiaochun Cao, Qingming Huang, Yuan Yao

HIGHLIGHT: In this paper, we construct a deep SVP prediction model which not only leads to better detection of annotation outliers but also enables learning with extremely sparse annotations.

TITLE: Transferable AutoML by Model Sharing Over Grouped Datasets

http://openaccess.thecvf.com/content_CVPR_2019/html/Xue_Transferable_AutoML_by_Model_Sharing_Over_Grouped_Datasets_CVPR_2019_paper.html

AUTHORS: Chao Xue, Junchi Yan, Rong Yan, Stephen M. Chu, Yonggang Hu, Yonghua Lin

HIGHLIGHT: This paper presents a so-called transferable AutoML approach that leverages previously trained models to speed up the search process for new tasks and datasets.

TITLE: Learning Not to Learn: Training Deep Neural Networks With Biased Data

http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Learning_Not_to_Learn_Training_Deep_Neural_Networks_With_Biased_CVPR_2019_paper.html

AUTHORS: Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, Junmo Kim

HIGHLIGHT: We propose a novel regularization algorithm to train deep neural networks, in which data at training time is severely biased.

TITLE: IRLAS: Inverse Reinforcement Learning for Architecture Search

http://openaccess.thecvf.com/content_CVPR_2019/html/Guo_IRLAS_Inverse_Reinforcement_Learning_for_Architecture_Search_CVPR_2019_paper.html

AUTHORS: Minghao Guo, Zhao Zhong, Wei Wu, Dahua Lin, Junjie Yan

HIGHLIGHT: In this paper, we propose an inverse reinforcement learning method for architecture search (IRLAS), which trains an agent to learn to search network structures that are topologically inspired by human-designed network.

TITLE: Learning for Single-Shot Confidence Calibration in Deep Neural Networks Through Stochastic Inferences

http://openaccess.thecvf.com/content_CVPR_2019/html/Seo_Learning_for_Single-Shot_Confidence_Calibration_in_Deep_Neural_Networks_Through_CVPR_2019_paper.html

AUTHORS: Seonguk Seo, Paul Hongsuck Seo, Bohyung Han

HIGHLIGHT: We propose a generic framework to calibrate accuracy and confidence of a prediction in deep neural networks through stochastic inferences.

TITLE: Attention-Based Adaptive Selection of Operations for Image Restoration in the Presence of Unknown Combined Distortions

http://openaccess.thecvf.com/content_CVPR_2019/html/Suganuma_Attention-Based_Adaptive_Selection_of_Operations_for_Image_Restoration_in_the_CVPR_2019_paper.html

AUTHORS: Masanori Suganuma, Xing Liu, Takayuki Okatani

HIGHLIGHT: For this purpose, we propose a simple yet effective layer architecture of neural networks.

TITLE: Fully Learnable Group Convolution for Acceleration of Deep Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Fully_Learnable_Group_Convolution_for_Acceleration_of_Deep_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Xijun Wang, Meina Kan, Shiguang Shan, Xilin Chen

HIGHLIGHT: To reduce the high computational and memory cost, in this work, we propose a fully learnable group convolution module (FLGC for short) which is quite efficient and can be embedded into any deep neural networks for acceleration.

TITLE: EIGEN: Ecologically-Inspired GENetic Approach for Neural Network Structure Searching From Scratch

http://openaccess.thecvf.com/content_CVPR_2019/html/Ren_EIGEN_Ecologically-Inspired_GENetic_Approach_for_Neural_Network_Structure_Searching_From_CVPR_2019_paper.html

AUTHORS: Jian Ren, Zhe Li, Jianchao Yang, Ning Xu, Tianbao Yang, David J. Foran

HIGHLIGHT: In this paper, we propose an Ecologically-Inspired GENetic (EIGEN) approach that uses the concept of succession, extinction, mimicry, and gene duplication to search neural network structure from scratch with poorly initialized simple network and few constraints forced during the evolution, as we assume no prior knowledge about the task domain.

TITLE: Deep Incremental Hashing Network for Efficient Image Retrieval

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Deep_Incremental_Hashing_Network_for_Efficient_Image_Retrieval_CVPR_2019_paper.html

AUTHORS: Dayan Wu, Qi Dai, Jing Liu, Bo Li, Weiping Wang

HIGHLIGHT: In this paper, we propose a novel deep hashing framework, called Deep Incremental Hashing Network (DIHN), for learning hash codes in an incremental manner.

TITLE: Robustness via Curvature Regularization, and Vice Versa

http://openaccess.thecvf.com/content_CVPR_2019/html/Moosavi-Dezfooli_Robustness_via_Curvature_Regularization_and_Vice_Versa_CVPR_2019_paper.html

AUTHORS: Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, Pascal Frossard

HIGHLIGHT: In this paper, we investigate the effect of adversarial training on the geometry of the classification landscape and decision boundaries.

TITLE: SparseFool: A Few Pixels Make a Big Difference

http://openaccess.thecvf.com/content_CVPR_2019/html/Modas_SparseFool_A_Few_Pixels_Make_a_Big_Difference_CVPR_2019_paper.html

AUTHORS: Apostolos Modas, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard

HIGHLIGHT: In this paper, we exploit the low mean curvature of the decision boundary, and propose SparseFool, a geometry inspired sparse attack that controls the sparsity of the perturbations.

TITLE: Interpretable and Fine-Grained Visual Explanations for Convolutional Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Wagner_Interpretable_and_Fine-Grained_Visual_Explanations_for_Convolutional_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Jorg Wagner, Jan Mathias Kohler, Tobias Gindele, Leon Hetzel, Jakob Thaddaus Wiedemer, Sven Behnke

HIGHLIGHT: In this work, we propose a post-hoc, optimization based visual explanation method, which highlights the evidence in the input image for a specific prediction.

TITLE: Structured Pruning of Neural Networks With Budget-Aware Regularization

http://openaccess.thecvf.com/content_CVPR_2019/html/Lemaire_Structured_Pruning_of_Neural_Networks_With_Budget-Aware_Regularization_CVPR_2019_paper.html

AUTHORS: Carl Lemaire, Andrew Achkar, Pierre-Marc Jodoin

HIGHLIGHT: To overcome this, we introduce a budgeted regularized pruning framework for deep CNNs.

TITLE: MBS: MacroblocK Scaling for CNN Model Reduction

http://openaccess.thecvf.com/content_CVPR_2019/html/Lin_MBS_MacroblocK_Scaling_for_CNN_Model_Reduction_CVPR_2019_paper.html

AUTHORS: Yu-Hsun Lin, Chun-Nan Chou, Edward Y. Chang

HIGHLIGHT: In this paper we propose the macroblock scaling (MBS) algorithm, which can be applied to various CNN architectures to reduce their model size.

TITLE: Fast Neural Architecture Search of Compact Semantic Segmentation Models via Auxiliary Cells

http://openaccess.thecvf.com/content_CVPR_2019/html/Nekrasov_Fast_Neural_Architecture_Search_of_Compact_Semantic_Segmentation_Models_via_CVPR_2019_paper.html

AUTHORS: Vladimir Nekrasov, Hao Chen, Chunhua Shen, Ian Reid

HIGHLIGHT: In this work, we are particularly interested in searching for high-performance compact segmentation architectures, able to run in real-time using limited resources.

TITLE: Generating 3D Adversarial Point Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Xiang_Generating_3D_Adversarial_Point_Clouds_CVPR_2019_paper.html
AUTHORS: Chong Xiang, Charles R. Qi, Bo Li
HIGHLIGHT: In this work, we propose several novel algorithms to craft adversarial point clouds against PointNet, a widely used deep neural network for point cloud processing.

TITLE: Partial Order Pruning: For Best Speed/Accuracy Trade-Off in Neural Architecture Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Partial_Order_Pruning_For_Best_SpeedAccuracy_Trade-Off_in_Neural_Architecture_CVPR_2019_paper.html
AUTHORS: Xin Li, Yiming Zhou, Zheng Pan, Jiashi Feng
HIGHLIGHT: In this work, we propose an algorithm that can offer better speed/accuracy trade-off of searched networks, which is termed "Partial Order Pruning".

TITLE: Memory in Memory: A Predictive Neural Network for Learning Higher-Order Non-Stationarity From Spatiotemporal Dynamics
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Memory_in_Memory_A_Predictive_Neural_Network_for_Learning_Higher-Order_CVPR_2019_paper.html
AUTHORS: Yunbo Wang, Jianjin Zhang, Hongyu Zhu, Mingsheng Long, Jianmin Wang, Philip S. Yu
HIGHLIGHT: We believe that the general idea of this work can be potentially applied to other time-series forecasting tasks.

TITLE: Variational Information Distillation for Knowledge Transfer
http://openaccess.thecvf.com/content_CVPR_2019/html/Ahn_Variational_Information_Distillation_for_Knowledge_Transfer_CVPR_2019_paper.html
AUTHORS: Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D. Lawrence, Zhenwen Dai
HIGHLIGHT: We propose an information-theoretic framework for knowledge transfer which formulates knowledge transfer as maximizing the mutual information between the teacher and the student networks.

TITLE: You Look Twice: GaterNet for Dynamic Filter Selection in CNNs
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_You_Look_Twice_GaterNet_for_Dynamic_Filter_Selection_in_CNNs_CVPR_2019_paper.html
AUTHORS: Zhouong Chen, Yang Li, Samy Bengio, Si Si
HIGHLIGHT: In this paper, we investigate input-dependent dynamic filter selection in deep convolutional neural networks (CNNs).

TITLE: SpherePHD: Applying CNNs on a Spherical PolyHeDron Representation of 360deg Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Lee_SpherePHD_Applying_CNNs_on_a_Spherical_PolyHeDron_Representation_of_360deg_CVPR_2019_paper.html
AUTHORS: Yeonkun Lee, Jaeseok Jeong, Jongseob Yun, Wonjune Cho, Kuk-Jin Yoon
HIGHLIGHT: This paper presents a novel method to resolve such problems of applying CNNs to omni-directional images.

TITLE: ESPNetv2: A Light-Weight, Power Efficient, and General Purpose Convolutional Neural Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Mehta_ESPNetv2_A_Light-Weight_Power_Efficient_and_General_Purpose_Convolutional_Neural_CVPR_2019_paper.html
AUTHORS: Sachin Mehta, Mohammad Rastegari, Linda Shapiro, Hannaneh Hajishirzi
HIGHLIGHT: We introduce a light-weight, power efficient, and general purpose convolutional neural network, ESPNetv2, for modeling visual and sequential data.

TITLE: Assisted Excitation of Activations: A Learning Technique to Improve Object Detectors
http://openaccess.thecvf.com/content_CVPR_2019/html/Derakhshani_Assisted_Excitation_of_Activations_A_Learning_Technique_to_Improve_Object_CVPR_2019_paper.html
AUTHORS: Mohammad Mahdi Derakhshani, Saeed Masoudnia, Amir Hossein Shaker, Omid Mersa, Mohammad Amin Sadeghi, Mohammad Rastegari, Babak N. Araabi
HIGHLIGHT: We present a simple yet effective learning technique that significantly improves mAP of YOLO object detectors without compromising their speed.

TITLE: Exploiting Edge Features for Graph Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Gong_Exploiting_Edge_Features_for_Graph_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Liyu Gong, Qiang Cheng
HIGHLIGHT: In this paper, we build a new framework for a family of new graph neural network models that can more sufficiently exploit edge features, including those of undirected or multi-dimensional edges.

TITLE: Propagation Mechanism for Deep and Wide Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_Propagation_Mechanism_for_Deep_and_Wide_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Dejiang Xu, Mong Li Lee, Wynne Hsu

HIGHLIGHT: In this paper, we propose a new propagation mechanism called channel-wise addition (cAdd) to deal with the vanishing gradients problem without sacrificing the complexity of the learned features.

TITLE: Catastrophic Child's Play: Easy to Perform, Hard to Defend Adversarial Attacks
http://openaccess.thecvf.com/content_CVPR_2019/html/Ho_Catastrophic_Childs_Play_Easy_to_Perform_Hard_to_Defend_Adversarial_CVPR_2019_paper.html

AUTHORS: Chih-Hui Ho, Brandon Leung, Erik Sandstrom, Yen Chang, Nuno Vasconcelos

HIGHLIGHT: A framework for the study of such attacks is proposed, using real world object manipulations.

TITLE: Embedding Complementary Deep Networks for Image Classification
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Embedding_Complementary_Deep_Networks_for_Image_Classification_CVPR_2019_paper.html

AUTHORS: Qiuyu Chen, Wei Zhang, Jun Yu, Jianping Fan

HIGHLIGHT: In this paper, a deep embedding algorithm is developed to achieve higher accuracy rates on large-scale image classification.

TITLE: Practical Full Resolution Learned Lossless Image Compression
http://openaccess.thecvf.com/content_CVPR_2019/html/Mentzer_Practical_Full_Resolution_Learned_Lossless_Image_Compression_CVPR_2019_paper.html

AUTHORS: Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, Luc Van Gool

HIGHLIGHT: We propose the first practical learned lossless image compression system, L3C, and show that it outperforms the popular engineered codecs, PNG, WebP and JPEG 2000.

TITLE: Image-To-Image Translation via Group-Wise Deep Whitening-And-Coloring Transformation
http://openaccess.thecvf.com/content_CVPR_2019/html/Cho_Image-To-Image_Translation_via_Group-Wise_Deep_Whitening-And-Coloring_Transformation_CVPR_2019_paper.html

AUTHORS: Wonwoong Cho, Sungha Choi, David Keetae Park, Inkyu Shin, Jaegul Choo

HIGHLIGHT: In response, this paper proposes an end-to-end approach tailored for image translation that efficiently approximates this transformation with our novel regularization methods.

TITLE: Max-Sliced Wasserstein Distance and Its Use for GANs
http://openaccess.thecvf.com/content_CVPR_2019/html/Deshpande_Max-Sliced_Wasserstein_Distance_and_Its_Use_for_GANs_CVPR_2019_paper.html

AUTHORS: Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, Alexander G. Schwing

HIGHLIGHT: Max-Sliced Wasserstein Distance and Its Use for GANs

TITLE: Meta-Learning With Differentiable Convex Optimization
http://openaccess.thecvf.com/content_CVPR_2019/html/Lee_Meta-Learning_With_Differentiable_Convex_Optimization_CVPR_2019_paper.html

AUTHORS: Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, Stefano Soatto

HIGHLIGHT: Our objective is to learn feature embeddings that generalize well under a linear classification rule for novel categories.

TITLE: RePr: Improved Training of Convolutional Filters
http://openaccess.thecvf.com/content_CVPR_2019/html/Prakash_RePr_Improved_Training_of_Convolutional_Filters_CVPR_2019_paper.html

AUTHORS: Aaditya Prakash, James Storer, Dinei Florencio, Cha Zhang

HIGHLIGHT: Our method is applicable both to vanilla convolutional networks and more complex modern architectures, and improves the performance across a variety of tasks, especially when applied to smaller networks.

TITLE: Tangent-Normal Adversarial Regularization for Semi-Supervised Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Yu_Tangent-Normal_Adversarial_Regularization_for_Semi-Supervised_Learning_CVPR_2019_paper.html

AUTHORS: Bing Yu, Jingfeng Wu, Jinwen Ma, Zhanxing Zhu

HIGHLIGHT: In this work, we propose tangent-normal adversarial regularization (TNAR) as an extension of VAT by taking the data manifold into consideration.

TITLE: Auto-Encoding Scene Graphs for Image Captioning

http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Auto-Encoding_Scene_Graphs_for_Image_Captioning_CVPR_2019_paper.html

AUTHORS: Xu Yang, Kaihua Tang, Hanwang Zhang, Jianfei Cai

HIGHLIGHT: We propose Scene Graph Auto-Encoder (SGAE) that incorporates the language inductive bias into the encoder-decoder image captioning framework for more human-like captions.

TITLE: Fast, Diverse and Accurate Image Captioning Guided by Part-Of-Speech

http://openaccess.thecvf.com/content_CVPR_2019/html/Deshpande_Fast_Diverse_and_Accurate_Image_Captioning_Guided_by_Part-Of-Speech_CVPR_2019_paper.html

AUTHORS: Aditya Deshpande, Jyoti Aneja, Liwei Wang, Alexander G. Schwing, David Forsyth

HIGHLIGHT: We use part-of-speech as summaries, since our summary should drive caption generation.

TITLE: Attention Branch Network: Learning of Attention Mechanism for Visual Explanation

http://openaccess.thecvf.com/content_CVPR_2019/html/Fukui_Attention_Branch_Network_Learning_of_Attention_Mechanism_for_Visual_Explanation_CVPR_2019_paper.html

AUTHORS: Hiroshi Fukui, Tsubasa Hirakawa, Takayoshi Yamashita, Hironobu Fujiyoshi

HIGHLIGHT: In this paper, we focus on the attention map for visual explanation, which represents a high response value as the attention location in image recognition.

TITLE: Cascaded Projection: End-To-End Network Compression and Acceleration

http://openaccess.thecvf.com/content_CVPR_2019/html/Minnehan_Cascaded_Projection_End-To-End_Network_Compression_and_Acceleration_CVPR_2019_paper.html

AUTHORS: Breton Minnehan, Andreas Savakis

HIGHLIGHT: We propose a data-driven approach for deep convolutional neural network compression that achieves high accuracy with high throughput and low memory requirements.

TITLE: DeepCaps: Going Deeper With Capsule Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Rajasegaran_DeepCaps_Going_Deeper_With_Capsule_Networks_CVPR_2019_paper.html

AUTHORS: Jathushan Rajasegaran, Vinoj Jayasundara, Sandaru Jayasekara, Hirunima Jayasekara, Suranga Seneviratne, Ranga Rodrigo

HIGHLIGHT: Drawing intuition from the success achieved by Convolutional Neural Networks (CNNs) by going deeper, we introduce DeepCaps, a deep capsule network architecture which uses a novel 3D convolution based dynamic routing algorithm.

TITLE: FBNet: Hardware-Aware Efficient ConvNet Design via Differentiable Neural Architecture Search

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_FBNet_Hardware-Aware_Efficient_ConvNet_Design_via_Differentiable_Neural_Architecture_Search_CVPR_2019_paper.html

AUTHORS: Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, Kurt Keutzer

HIGHLIGHT: To address these, we propose a differentiable neural architecture search (DNAS) framework that uses gradient-based methods to optimize ConvNet architectures, avoiding enumerating and training individual architectures separately as in previous methods.

TITLE: APDrawingGAN: Generating Artistic Portrait Drawings From Face Photos With Hierarchical GANs

http://openaccess.thecvf.com/content_CVPR_2019/html/Yi_APDrawingGAN_Generating_Artistic_Portrait_Drawings_From_Face_Photos_With_Hierarchical_CVPR_2019_paper.html

AUTHORS: Ran Yi, Yong-Jin Liu, Yu-Kun Lai, Paul L. Rosin

HIGHLIGHT: To address these challenges, we propose APDrawingGAN, a novel GAN based architecture that builds upon hierarchical generators and discriminators combining both a global network (for images as a whole) and local networks (for individual facial regions).

TITLE: Constrained Generative Adversarial Networks for Interactive Image Generation

http://openaccess.thecvf.com/content_CVPR_2019/html/Heim_Constrained_Generative_Adversarial_Networks_for_Interactive_Image_Generation_CVPR_2019_paper.html

AUTHORS: Eric Heim

HIGHLIGHT: In this work we develop a novel GAN framework that allows humans to be "in-the-loop" of the image generation process.

TITLE: WarpGAN: Automatic Caricature Generation
http://openaccess.thecvf.com/content_CVPR_2019/html/Shi_WarpGAN_Automatic_Caricature_Generation_CVPR_2019_paper.html
AUTHORS: Yichun Shi, Debayan Deb, Anil K. Jain
HIGHLIGHT: We propose, WarpGAN, a fully automatic network that can generate caricatures given an input face photo.

TITLE: Explainability Methods for Graph Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Pope_Explainability_Methods_for_Graph_Convolutional_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Phillip E. Pope, Soheil Kolouri, Mohammad Rostami, Charles E. Martin, Heiko Hoffmann
HIGHLIGHT: In this paper, we introduce explainability methods for GCNNs.

TITLE: A Generative Adversarial Density Estimator
http://openaccess.thecvf.com/content_CVPR_2019/html/Abbasnejad_A_Generative_Adversarial_Density_Estimator_CVPR_2019_paper.html
AUTHORS: M. Ehsan Abbasnejad, Qinfeng Shi, Anton van den Hengel, Lingqiao Liu
HIGHLIGHT: We propose a Generative Adversarial Density Estimator, a density estimation approach that bridges the gap between the two.

TITLE: SoDeep: A Sorting Deep Net to Learn Ranking Loss Surrogates
http://openaccess.thecvf.com/content_CVPR_2019/html/Engilberge_SoDeep_A_Sorting_Deep_Net_to_Learn_Ranking_Loss_Surrogates_CVPR_2019_paper.html
AUTHORS: Martin Engilberge, Louis Chevallier, Patrick Perez, Matthieu Cord
HIGHLIGHT: In the present work, we introduce a new method to learn approximations of such non-differentiable objective functions.

TITLE: Pixel-Adaptive Convolutional Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Su_Pixel-Adaptive_Convolutional_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Hang Su, Varun Jampani, Deqing Sun, Orazio Gallo, Erik Learned-Miller, Jan Kautz
HIGHLIGHT: We propose a pixel-adaptive convolution (PAC) operation, a simple yet effective modification of standard convolutions, in which the filter weights are multiplied with a spatially varying kernel that depends on learnable, local pixel features.

TITLE: Single-Frame Regularization for Temporally Stable CNNs
http://openaccess.thecvf.com/content_CVPR_2019/html/Eilertsen_Single-Frame_Regularization_for_Temporally_Stable_CNNs_CVPR_2019_paper.html
AUTHORS: Gabriel Eilertsen, Rafal K. Mantiuk, Jonas Unger
HIGHLIGHT: We take a different approach to the problem, posing temporal stability as a regularization of the cost function.

TITLE: An End-To-End Network for Generating Social Relationship Graphs
http://openaccess.thecvf.com/content_CVPR_2019/html/Goel_An_End-To-End_Network_for_Generating_Social_Relationship_Graphs_CVPR_2019_paper.html
AUTHORS: Arushi Goel, Keng Teck Ma, Cheston Tan
HIGHLIGHT: We introduce a novel end-to-end-trainable neural network that is capable of generating a Social Relationship Graph - a structured, unified representation of social relationships and attributes - from a given input image.

TITLE: Meta-Learning Convolutional Neural Architectures for Multi-Target Concrete Defect Classification With the COConcrete Defect BRidge IMage Dataset
http://openaccess.thecvf.com/content_CVPR_2019/html/Mundt_Meta-Learning_Convolutional_Neural_Architectures_for_Multi-Target_Concrete_Defect_Classification_With_CVPR_2019_paper.html
AUTHORS: Martin Mundt, Sagnik Majumder, Sreenivas Murali, Panagiotis Panetsos, Visvanathan Ramesh
HIGHLIGHT: In this work we introduce the novel COConcrete Defect BRidge IMage dataset (CODEBRIM) for multi-target classification of five commonly appearing concrete defects.

TITLE: ECC: Platform-Independent Energy-Constrained Deep Neural Network Compression via a Bilinear Regression Model
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_ECC_Platform-Independent_Energy-Constrained_Deep_Neural_Network_Compression_via_a_Bilinear_CVPR_2019_paper.html
AUTHORS: Haichuan Yang, Yuhao Zhu, Ji Liu
HIGHLIGHT: This paper proposes ECC, a framework that compresses DNNs to meet a given energy constraint while minimizing accuracy loss.

TITLE: SeerNet: Predicting Convolutional Neural Network Feature-Map Sparsity Through Low-Bit Quantization
http://openaccess.thecvf.com/content_CVPR_2019/html/Cao_SeerNet_Predicting_Convolutional_Neural_Network_Feature-Map_Sparsity_Through_Low-Bit_Quantization_CVPR_2019_paper.html
AUTHORS: Shijie Cao, Lingxiao Ma, Wencong Xiao, Chen Zhang, Yunxin Liu, Lintao Zhang, Lanshun Nie, Zhi Yang
HIGHLIGHT: In this paper we present a novel and general method to accelerate convolutional neural network (CNN) inference by taking advantage of feature map sparsity.

TITLE: Defending Against Adversarial Attacks by Randomized Diversification
http://openaccess.thecvf.com/content_CVPR_2019/html/Taran_Defending_Against_Adversarial_Attacks_by_Randomized_Diversification_CVPR_2019_paper.html
AUTHORS: Olga Taran, Shideh Rezaeifar, Taras Holotyak, Slava Voloshynovskiy
HIGHLIGHT: In this paper, we propose a randomized diversification as a defense strategy.

TITLE: Rob-GAN: Generator, Discriminator, and Adversarial Attacker
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Rob-GAN_Generator_Discriminator_and_Adversarial_Attacker_CVPR_2019_paper.html
AUTHORS: Xuanqing Liu, Cho-Jui Hsieh
HIGHLIGHT: Combining these two insights, we develop a framework called Rob-GAN to jointly optimize generator and discriminator in the presence of adversarial attacks---the generator generates fake images to fool discriminator; the adversarial attacker perturbs real images to fool discriminator, and the discriminator wants to minimize loss under fake and adversarial images.

TITLE: Learning From Noisy Labels by Regularized Estimation of Annotator Confusion
http://openaccess.thecvf.com/content_CVPR_2019/html/Tanno_Learning_From_Noisy_Labels_by_Regularized_Estimation_of_Annotator_Confusion_CVPR_2019_paper.html
AUTHORS: Ryutaro Tanno, Ardavan Saeedi, Swami Sankaranarayanan, Daniel C. Alexander, Nathan Silberman
HIGHLIGHT: In this work, we present a method for simultaneously learning the individual annotator model and the underlying true label distribution, using only noisy observations.

TITLE: Task-Free Continual Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Aljundi_Task-Free_Continual_Learning_CVPR_2019_paper.html
AUTHORS: Rahaf Aljundi, Klaas Kelchtermans, Tinne Tuytelaars
HIGHLIGHT: Therefore we investigate how to transform continual learning to an online setup.

TITLE: Importance Estimation for Neural Network Pruning
http://openaccess.thecvf.com/content_CVPR_2019/html/Molchanov_Importance_Estimation_for_Neural_Network_Pruning_CVPR_2019_paper.html
AUTHORS: Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, Jan Kautz
HIGHLIGHT: We propose a novel method that estimates the contribution of a neuron (filter) to the final loss and iteratively removes those with smaller scores.

TITLE: Detecting Overfitting of Deep Generative Networks via Latent Recovery
http://openaccess.thecvf.com/content_CVPR_2019/html/Webster_Detecting_Overfitting_of_Deep_Generative_Networks_via_Latent_Recovery_CVPR_2019_paper.html
AUTHORS: Ryan Webster, Julien Rabin, Loic Simon, Frederic Jurie
HIGHLIGHT: We address this question by i) showing how simple losses are highly effective at reconstructing images for deep generators ii) analyzing the statistics of reconstruction errors for training versus validation images.

TITLE: Coloring With Limited Data: Few-Shot Colorization via Memory Augmented Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Yoo_Coloring_With_Limited_Data_Few-Shot_Colorization_via_Memory_Augmented_Networks_CVPR_2019_paper.html
AUTHORS: Seungjoo Yoo, Hyojin Bahng, Sunghyo Chung, Junsoo Lee, Jaehyuk Chang, Jaegul Choo
HIGHLIGHT: To tackle this issue, we present a novel memory-augmented colorization model MemoPainter that can produce high-quality colorization with limited data.

TITLE: Characterizing and Avoiding Negative Transfer
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Characterizing_and_Avoiding_Negative_Transfer_CVPR_2019_paper.html
AUTHORS: Zirui Wang, Zihang Dai, Barnabas Poczos, Jaime Carbonell
HIGHLIGHT: This paper proposes a formal definition of negative transfer and analyzes three important aspects thereof.

TITLE: Building Efficient Deep Neural Networks With Unitary Group Convolutions

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Building_Efficient_Deep_Neural_Networks_With_Unitary_Group_Convolutions_CVPR_2019_paper.html

AUTHORS: Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Christopher De Sa, Zhiru Zhang

HIGHLIGHT: We propose unitary group convolutions (UGConvs), a building block for CNNs which compose a group convolution with unitary transforms in feature space to learn a richer set of representations than group convolution alone.

TITLE: Semi-Supervised Learning With Graph Learning-Convolutional Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Jiang_Semi-Supervised_Learning_With_Graph_Learning-Convolutional_Networks_CVPR_2019_paper.html

AUTHORS: Bo Jiang, Ziyang Zhang, Doudou Lin, Jin Tang, Bin Luo

HIGHLIGHT: In this paper, we propose a novel Graph Learning-Convolutional Network (GLCN) for graph data representation and semi-supervised learning.

TITLE: Learning to Remember: A Synaptic Plasticity Driven Framework for Continual Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Ostapenko_Learning_to_Remember_A_Synaptic_Plasticity_Driven_Framework_for_Continual_CVPR_2019_paper.html

AUTHORS: Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jahnichen, Moin Nabi

HIGHLIGHT: In order to tackle these challenges, we introduce Dynamic Generative Memory (DGM) - synaptic plasticity driven framework for continual learning.

TITLE: AIRD: Adversarial Learning Framework for Image Repurposing Detection

http://openaccess.thecvf.com/content_CVPR_2019/html/Jaiswal_AIRD_Adversarial_Learning_Framework_for_Image_Repurposing_Detection_CVPR_2019_paper.html

AUTHORS: Ayush Jaiswal, Yue Wu, Wael AbdAlmageed, Iacopo Masi, Premkumar Natarajan

HIGHLIGHT: In this paper, we present a novel method for image repurposing detection that is based on the real-world adversarial interplay between a bad actor who repurposes images with counterfeit metadata and a watchdog who verifies the semantic consistency between images and their accompanying metadata, where both players have access to a reference dataset of verified content, which they can use to achieve their goals.

TITLE: A Kernelized Manifold Mapping to Diminish the Effect of Adversarial Perturbations

http://openaccess.thecvf.com/content_CVPR_2019/html/Taghanaki_A_Kernelized_Manifold_Mapping_to_Diminish_the_Effect_of_Adversarial_CVPR_2019_paper.html

AUTHORS: Saeid Asgari Taghanaki, Kumar Abhishek, Shekoofeh Azizi, Ghassan Hamarneh

HIGHLIGHT: To tackle this problem, we propose a non-linear radial basis convolutional feature mapping by learning a Mahalanobis-like distance function.

TITLE: Trust Region Based Adversarial Attack on Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Yao_Trust_Region_Based_Adversarial_Attack_on_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Zhewei Yao, Amir Gholami, Peng Xu, Kurt Keutzer, Michael W. Mahoney

HIGHLIGHT: To address this problem, we present a new family of trust region based adversarial attacks, with the goal of computing adversarial perturbations efficiently.

TITLE: PEPSI: Fast Image Inpainting With Parallel Decoding Network

http://openaccess.thecvf.com/content_CVPR_2019/html/Sagong_PEPSI_Fast_Image_Inpainting_With_Parallel_Decoding_Network_CVPR_2019_paper.html

AUTHORS: Min-cheol Sagong, Yong-goo Shin, Seung-wook Kim, Seung Park, Sung-jea Ko

HIGHLIGHT: To solve this problem, in this paper, we present a novel network structure, called PEPSI: parallel extended-decoder path for semantic inpainting.

TITLE: Model-Blind Video Denoising via Frame-To-Frame Training

http://openaccess.thecvf.com/content_CVPR_2019/html/Ehret_Model-Blind_Video_Denoising_via_Frame-To-Frame_Training_CVPR_2019_paper.html

AUTHORS: Thibaud Ehret, Axel Davy, Jean-Michel Morel, Gabriele Facciolo, Pablo Arias

HIGHLIGHT: In this paper we propose a fully blind video denoising method, with two versions off-line and on-line.

TITLE: End-To-End Efficient Representation Learning via Cascading Combinatorial Optimization

http://openaccess.thecvf.com/content_CVPR_2019/html/Jeong_End-To-End_Efficient_Representation_Learning_via_Cascading_Combinatorial_Optimization_CVPR_2019_paper.html

AUTHORS: Yeonwoo Jeong, Yoonsung Kim, Hyun Oh Song

HIGHLIGHT: We develop hierarchically quantized efficient embedding representations for similarity-based search and show that this representation provides not only the state of the art performance on the search accuracy but also provides several orders of speed up during inference.

TITLE: Sim-Real Joint Reinforcement Transfer for 3D Indoor Navigation

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhu_Sim-Real_Joint_Reinforcement_Transfer_for_3D_Indoor_Navigation_CVPR_2019_paper.html

AUTHORS: Fengda Zhu, Linchao Zhu, Yi Yang

HIGHLIGHT: Specifically, our method employs an adversarial feature adaptation model for visual representation transfer and a policy mimic strategy for policy behavior imitation.

TITLE: ChamNet: Towards Efficient Network Design Through Platform-Aware Model Adaptation

http://openaccess.thecvf.com/content_CVPR_2019/html/Dai_ChamNet_Towards_Efficient_Network_Design_Through_Platform-Aware_Model_Adaptation_CVPR_2019_paper.html

AUTHORS: Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, Peter Vajda, Matt Uyttendaele, Niraj K. Jha

HIGHLIGHT: We formulate platform-aware NN architecture search in an optimization framework and propose a novel algorithm to search for optimal architectures aided by efficient accuracy and resource (latency and/or energy) predictors.

TITLE: Regularizing Activation Distribution for Training Binarized Deep Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Ding_Regularizing_Activation_Distribution_for_Training_Binarized_Deep_Networks_CVPR_2019_paper.html

AUTHORS: Ruizhou Ding, Ting-Wu Chin, Zeye Liu, Diana Marculescu

HIGHLIGHT: In this paper, we propose to use distribution loss to explicitly regularize the activation flow, and develop a framework to systematically formulate the loss.

TITLE: Robustness Verification of Classification Deep Neural Networks via Linear Programming

http://openaccess.thecvf.com/content_CVPR_2019/html/Lin_Robustness_Verification_of_Classification_Deep_Neural_Networks_via_Linear_Programming_CVPR_2019_paper.html

AUTHORS: Wang Lin, Zhengfeng Yang, Xin Chen, Qingye Zhao, Xiangkun Li, Zhiming Liu, Jifeng He

HIGHLIGHT: In this paper, we develop a novel method for robustness verification of CDNNs with sigmoid activation functions.

TITLE: Additive Adversarial Learning for Unbiased Authentication

http://openaccess.thecvf.com/content_CVPR_2019/html/Liang_Additive_Adversarial_Learning_for_Unbiased_Authentication_CVPR_2019_paper.html

AUTHORS: Jian Liang, Yuren Cao, Chenbin Zhang, Shiyu Chang, Kun Bai, Zenglin Xu

HIGHLIGHT: To address this issue, we propose a novel two-stage method that disentangles the class/identity from domain-differences, and we consider multiple types of domain-difference.

TITLE: Simultaneously Optimizing Weight and Quantizer of Ternary Neural Network Using Truncated Gaussian Approximation

http://openaccess.thecvf.com/content_CVPR_2019/html/He_Simultaneously_Optimizing_Weight_and_Quantizer_of_Ternary_Neural_Network_Using_CVPR_2019_paper.html

AUTHORS: Zhezhi He, Deliang Fan

HIGHLIGHT: In this work, we propose a novel ternarized neural network training method which simultaneously optimizes both weights and quantizer during training, differentiating from prior works.

TITLE: Adversarial Defense by Stratified Convolutional Sparse Coding

http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Adversarial_Defense_by_Stratified_Convolutional_Sparse_Coding_CVPR_2019_paper.html

AUTHORS: Bo Sun, Nian-Hsuan Tsai, Fangchen Liu, Ronald Yu, Hao Su

HIGHLIGHT: We propose an adversarial defense method that achieves state-of-the-art performance among attack-agnostic adversarial defense methods while also maintaining robustness to input resolution, scale of adversarial perturbation, and scale of dataset size.