

TITLE: DrivingStereo: A Large-Scale Dataset for Stereo Matching in Autonomous Driving Scenarios
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_DrivingStereo_A_Large-Scale_Dataset_for_Stereo_Matching_in_Autonomous_Driving_CVPR_2019_paper.html

AUTHORS: Guorun Yang, Xiao Song, Chaoqin Huang, Zhidong Deng, Jianping Shi, Bolei Zhou
HIGHLIGHT: In this paper, we construct a novel large-scale stereo dataset named DrivingStereo.

TITLE: PartNet: A Large-Scale Benchmark for Fine-Grained and Hierarchical Part-Level 3D Object Understanding
http://openaccess.thecvf.com/content_CVPR_2019/html/Mo_PartNet_A_Large-Scale_Benchmark_for_Fine-Grained_and_Hierarchical_Part-Level_3D_CVPR_2019_paper.html

AUTHORS: Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, Hao Su
HIGHLIGHT: We present PartNet: a consistent, large-scale dataset of 3D objects annotated with fine-grained, instance-level, and hierarchical 3D part information.

TITLE: A Dataset and Benchmark for Large-Scale Multi-Modal Face Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_A_Dataset_and_Benchmark_for_Large-Scale_Multi-Modal_Face_Anti-Spoofing_CVPR_2019_paper.html

AUTHORS: Shifeng Zhang, Xiaobo Wang, Ajian Liu, Chenxu Zhao, Jun Wan, Sergio Escalera, Hailin Shi, Zezheng Wang, Stan Z. Li
HIGHLIGHT: To facilitate face anti-spoofing research, we introduce a large-scale multi-modal dataset, namely CASIA-SURF, which is the largest publicly available dataset for face anti-spoofing in terms of both subjects and visual modalities.

TITLE: Unsupervised Learning of Consensus Maximization for 3D Vision Problems
http://openaccess.thecvf.com/content_CVPR_2019/html/Probst_Unsupervised_Learning_of_Consensus_Maximization_for_3D_Vision_Problems_CVPR_2019_paper.html

AUTHORS: Thomas Probst, Danda Pani Paudel, Ajad Chhatkuli, Luc Van Gool
HIGHLIGHT: In this paper, we propose for the first time an unsupervised learning framework for consensus maximization, in the context of solving 3D vision problems.

TITLE: VizWiz-Priv: A Dataset for Recognizing the Presence and Purpose of Private Visual Information in Images Taken by Blind People
http://openaccess.thecvf.com/content_CVPR_2019/html/Gurari_VizWiz-Priv_A_Dataset_for_Recognizing_the_Presence_and_Purpose_of_CVPR_2019_paper.html

AUTHORS: Danna Gurari, Qing Li, Chi Lin, Yinan Zhao, Anhong Guo, Abigale Stangl, Jeffrey P. Bigham
HIGHLIGHT: We introduce the first visual privacy dataset originating from people who are blind in order to better understand their privacy disclosures and to encourage the development of algorithms that can assist in preventing their unintended disclosures.

TITLE: InverseRenderNet: Learning Single Image Inverse Rendering
http://openaccess.thecvf.com/content_CVPR_2019/html/Yu_InverseRenderNet_Learning_Single_Image_Inverse_Rendering_CVPR_2019_paper.html

AUTHORS: Ye Yu, William A. P. Smith
HIGHLIGHT: Since the problem is ill-posed we introduce additional supervision: 1.

TITLE: A Variational Auto-Encoder Model for Stochastic Point Processes
http://openaccess.thecvf.com/content_CVPR_2019/html/Mehrasa_A_Variational_Auto-Encoder_Model_for_Stochastic_Point_Processes_CVPR_2019_paper.html

AUTHORS: Nazanin Mehrasa, Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, Greg Mori
HIGHLIGHT: We propose a novel probabilistic generative model for action sequences.

TITLE: Unifying Heterogeneous Classifiers With Distillation
http://openaccess.thecvf.com/content_CVPR_2019/html/Vongkulbhisal_Unifying_Heterogeneous_Classifiers_With_Distillation_CVPR_2019_paper.html

AUTHORS: Jayakorn Vongkulbhisal, Phongtharin Vinayavekhin, Marco Visentini-Scarzanella
HIGHLIGHT: In this paper, we study the problem of unifying knowledge from a set of classifiers with different architectures and target classes into a single classifier, given only a generic set of unlabelled data.

TITLE: Assessment of Faster R-CNN in Man-Machine Collaborative Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Deza_Assessment_of_Faster_R-CNN_in_Man-Machine_Collaborative_Search_CVPR_2019_paper.html

AUTHORS: Arturo Deza, Amit Surana, Miguel P. Eckstein
HIGHLIGHT: With the advent of modern expert systems driven by deep learning that supplement human experts (e.g. radiologists, dermatologists, surveillance scanners), we analyze how and when do such expert systems enhance human performance in a fine-grained small target visual search task.

TITLE: OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge
http://openaccess.thecvf.com/content_CVPR_2019/html/Marino_OK-VQA_A_Visual_Question_Answering_Benchmark_Requiring_External_Knowledge_CVPR_2019_paper.html
AUTHORS: Kenneth Marino, Mohammad Rastegari, Ali Farhadi, Roozbeh Mottaghi
HIGHLIGHT: In this paper, we address the task of knowledge-based visual question answering and provide a benchmark, called OK-VQA, where the image content is not sufficient to answer the questions, encouraging methods that rely on external knowledge resources.

TITLE: NDDR-CNN: Layerwise Feature Fusing in Multi-Task CNNs by Neural Discriminative Dimensionality Reduction
http://openaccess.thecvf.com/content_CVPR_2019/html/Gao_NDDR-CNN_Layerwise_Feature_Fusing_in_Multi-Task_CNNs_by_Neural_Discriminative_CVPR_2019_paper.html
AUTHORS: Yuan Gao, Jiayi Ma, Mingbo Zhao, Wei Liu, Alan L. Yuille
HIGHLIGHT: In this paper, we propose a novel Convolutional Neural Network (CNN) structure for general-purpose multi-task learning (MTL), which enables automatic feature fusing at every layer from different tasks.

TITLE: Spectral Metric for Dataset Complexity Assessment
http://openaccess.thecvf.com/content_CVPR_2019/html/Branchaud-Charron_Spectral_Metric_for_Dataset_Complexity_Assessment_CVPR_2019_paper.html
AUTHORS: Frederic Branchaud-Charron, Andrew Achkar, Pierre-Marc Jodoin
HIGHLIGHT: In this paper, we propose a new measure to gauge the complexity of image classification problems.

TITLE: ADCrowdNet: An Attention-Injective Deformable Convolutional Network for Crowd Understanding
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_ADCrowdNet_An_Attention-Injective_Deformable_Convolutional_Network_for_Crowd_Understanding_CVPR_2019_paper.html
AUTHORS: Ning Liu, Yongchao Long, Changqing Zou, Qun Niu, Li Pan, Hefeng Wu
HIGHLIGHT: We propose an attention-injective deformable convolutional network called ADCrowdNet for crowd understanding that can address the accuracy degradation problem of highly congested noisy scenes.

TITLE: VERI-Wild: A Large Dataset and a New Method for Vehicle Re-Identification in the Wild
http://openaccess.thecvf.com/content_CVPR_2019/html/Lou_VERI-Wild_A_Large_Dataset_and_a_New_Method_for_Vehicle_CVPR_2019_paper.html
AUTHORS: Yihang Lou, Yan Bai, Jun Liu, Shiqi Wang, Lingyu Duan
HIGHLIGHT: To promote the research of vehicle ReID in the wild, we collect a new dataset called VERI-Wild with the following distinct features: 1) The vehicle images are captured by a large surveillance system containing 174 cameras covering a large urban district (more than 200km²) 2) The camera network continuously captures vehicles for 24 hours in each day and lasts for 1 month.

TITLE: RAVEN: A Dataset for Relational and Analogical Visual REasoning
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_RAVEN_A_Dataset_for_Relational_and_Analogical_Visual_REasoning_CVPR_2019_paper.html
AUTHORS: Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, Song-Chun Zhu
HIGHLIGHT: In this work, we propose a new dataset, built in the context of Raven's Progressive Matrices (RPM) and aimed at lifting machine intelligence by associating vision with structural, relational, and analogical reasoning in a hierarchical representation.

TITLE: Surface Reconstruction From Normals: A Robust DGP-Based Discontinuity Preservation Approach
http://openaccess.thecvf.com/content_CVPR_2019/html/Xie_Surface_Reconstruction_From_Normals_A_Robust_DGP-Based_Discontinuity_Preservation_Approach_CVPR_2019_paper.html
AUTHORS: Wuyuan Xie, Miaohui Wang, Mingqiang Wei, Jianmin Jiang, Jing Qin
HIGHLIGHT: This paper introduces a robust approach to preserve the surface discontinuity in the discrete geometry way.

TITLE: DeepFashion2: A Versatile Benchmark for Detection, Pose Estimation, Segmentation and Re-Identification of Clothing Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Ge_DeepFashion2_A_Versatile_Benchmark_for_Detection_Pose_Estimation_Segmentation_and_CVPR_2019_paper.html
AUTHORS: Yuying Ge, Ruimao Zhang, Xiaogang Wang, Xiaoou Tang, Ping Luo
HIGHLIGHT: DeepFashion2: A Versatile Benchmark for Detection, Pose Estimation, Segmentation and Re-Identification of Clothing Images

TITLE: Jumping Manifolds: Geometry Aware Dense Non-Rigid Structure From Motion

http://openaccess.thecvf.com/content_CVPR_2019/html/Kumar_Jumping_Manifolds_Geometry_Aware_Dense_Non-Rigid_Structure_From_Motion_CVPR_2019_paper.html

AUTHORS: Suryansh Kumar

HIGHLIGHT: Given dense image feature correspondences of a non-rigidly moving object across multiple frames, this paper proposes an algorithm to estimate its 3D shape for each frame.

TITLE: LVIS: A Dataset for Large Vocabulary Instance Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Gupta_LVIS_A_Dataset_for_Large_Vocabulary_Instance_Segmentation_CVPR_2019_paper.html

AUTHORS: Agrim Gupta, Piotr Dollar, Ross Girshick

HIGHLIGHT: In this work, we introduce LVIS (pronounced 'el-vis'): a new dataset for Large Vocabulary Instance Segmentation.

TITLE: Fast Object Class Labelling via Speech

http://openaccess.thecvf.com/content_CVPR_2019/html/Gygli_Fast_Object_Class_Labelling_via_Speech_CVPR_2019_paper.html

AUTHORS: Michael Gygli, Vittorio Ferrari

HIGHLIGHT: Instead, we propose a new interface where classes are annotated via speech.

TITLE: LaSOT: A High-Quality Benchmark for Large-Scale Single Object Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Fan_LaSOT_A_High-Quality_Benchmark_for_Large-Scale_Single_Object_Tracking_CVPR_2019_paper.html

AUTHORS: Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, Haibin Ling

HIGHLIGHT: In this paper, we present LaSOT, a high-quality benchmark for Large-scale Single Object Tracking.

TITLE: Creative Flow+ Dataset

http://openaccess.thecvf.com/content_CVPR_2019/html/Shugrina_Creative_Flow_Dataset_CVPR_2019_paper.html

AUTHORS: Maria Shugrina, Ziheng Liang, Amlan Kar, Jiaman Li, Angad Singh, Karan Singh, Sanja Fidler

HIGHLIGHT: We present the Creative Flow+ Dataset, the first diverse multi-style artistic video dataset richly labeled with per-pixel optical flow, occlusions, correspondences, segmentation labels, normals, and depth.

TITLE: Weakly Supervised Open-Set Domain Adaptation by Dual-Domain Collaboration

http://openaccess.thecvf.com/content_CVPR_2019/html/Tan_Weakly_Supervised_Open-Set_Domain_Adaptation_by_Dual-Domain_Collaboration_CVPR_2019_paper.html

AUTHORS: Shuhan Tan, Jiening Jiao, Wei-Shi Zheng

HIGHLIGHT: To address this practical setting, we propose the Collaborative Distribution Alignment (CDA) method, which performs knowledge transfer bilaterally and works collaboratively to classify unlabeled data and identify outlier samples.

TITLE: A Neurobiological Evaluation Metric for Neural Network Model Search

http://openaccess.thecvf.com/content_CVPR_2019/html/Blanchard_A_Neurobiological_Evaluation_Metric_for_Neural_Network_Model_Search_CVPR_2019_paper.html

AUTHORS: Nathaniel Blanchard, Jeffery Kinnison, Brandon Richard Webster, Pouya Bashivan, Walter J. Scheirer

HIGHLIGHT: In this paper we introduce a human-model similarity (HMS) metric, which quantifies the similarity of human fMRI and network activation behavior.

TITLE: Iterative Projection and Matching: Finding Structure-Preserving Representatives and Its Application to Computer Vision

http://openaccess.thecvf.com/content_CVPR_2019/html/Zaeemzadeh_Iterative_Projection_and_Matching_Finding_Structure-Preserving_Representatives_and_Its_Application_CVPR_2019_paper.html

AUTHORS: Alireza Zaeemzadeh, Mohsen Joneidi, Nazanin Rahnavard, Mubarak Shah

HIGHLIGHT: This paper presents a fast and accurate data selection method, in which the selected samples are optimized to span the subspace of all data.

TITLE: Efficient Multi-Domain Learning by Covariance Normalization

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Efficient_Multi-Domain_Learning_by_Covariance_Normalization_CVPR_2019_paper.html

AUTHORS: Yunsheng Li, Nuno Vasconcelos

HIGHLIGHT: CovNorm is a data driven method of fairly simple implementation, requiring two principal component analyzes (PCA) and fine-tuning of a mini-adaptation layer.

TITLE: Predicting Visible Image Differences Under Varying Display Brightness and Viewing Distance

http://openaccess.thecvf.com/content_CVPR_2019/html/Ye_Predicting_Visible_Image_Differences_Under_Varying_Display_Brightness_and_Viewing_CVPR_2019_paper.html

AUTHORS: Nanyang Ye, Krzysztof Wolski, Rafal K. Mantiuk

HIGHLIGHT: In this paper, we propose a CNN-based visibility metric, which maintains the accuracy of deep network solutions and accounts for viewing conditions.

TITLE: A Bayesian Perspective on the Deep Image Prior

http://openaccess.thecvf.com/content_CVPR_2019/html/Cheng_A_Bayesian_Perspective_on_the_Deep_Image_Prior_CVPR_2019_paper.html

AUTHORS: Zezhou Cheng, Matheus Gadelha, Subhansu Maji, Daniel Sheldon

HIGHLIGHT: We show that the deep image prior is asymptotically equivalent to a stationary Gaussian process prior in the limit as the number of channels in each layer of the network goes to infinity, and derive the corresponding kernel.

TITLE: ApolloCar3D: A Large 3D Car Instance Understanding Benchmark for Autonomous Driving

http://openaccess.thecvf.com/content_CVPR_2019/html/Song_ApolloCar3D_A_Large_3D_Car_Instance_Understanding_Benchmark_for_Autonomous_CVPR_2019_paper.html

AUTHORS: Xibin Song, Peng Wang, Dingfu Zhou, Rui Zhu, Chenye Guan, Yuchao Dai, Hao Su, Hongdong Li, Ruiqiang Yang

HIGHLIGHT: In this paper, we contribute the first large scale database suitable for 3D car instance understanding - ApolloCar3D.

TITLE: Compressing Unknown Images With Product Quantizer for Efficient Zero-Shot Classification

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Compressing_Unknown_Images_With_Product_Quantizer_for_Efficient_Zero-Shot_Classification_CVPR_2019_paper.html

AUTHORS: Jin Li, Xuguang Lan, Yang Liu, Le Wang, Nanning Zheng

HIGHLIGHT: Based on this intuition, a Product Quantization Zero-Shot Learning (PQZSL) method is proposed to learn embeddings as well as quantizers to compress visual features into compact codes for Approximate NN (ANN) search.

TITLE: Self-Supervised Convolutional Subspace Clustering Network

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Self-Supervised_Convolutional_Subspace_Clustering_Network_CVPR_2019_paper.html

AUTHORS: Junjian Zhang, Chun-Guang Li, Chong You, Xianbiao Qi, Honggang Zhang, Jun Guo, Zhouchen Lin

HIGHLIGHT: To achieve simultaneous feature learning and subspace clustering, we propose an end-to-end trainable framework, called Self-Supervised Convolutional Subspace Clustering Network (S²ConvSCN), that combines a ConvNet module (for feature learning), a self-expression module (for subspace clustering) and a spectral clustering module (for self-supervision) into a joint optimization framework.

TITLE: Isospectralization, or How to Hear Shape, Style, and Correspondence

http://openaccess.thecvf.com/content_CVPR_2019/html/Cosmo_Isospectralization_or_How_to_Hear_Shape_Style_and_Correspondence_CVPR_2019_paper.html

AUTHORS: Luca Cosmo, Mikhail Panine, Arianna Rampini, Maks Ovsjanikov, Michael M. Bronstein, Emanuele Rodola

HIGHLIGHT: In this paper, we introduce a numerical procedure called isospectralization, consisting of deforming one shape to make its Laplacian spectrum match that of another.

TITLE: Speech2Face: Learning the Face Behind a Voice

http://openaccess.thecvf.com/content_CVPR_2019/html/Oh_Speech2Face_Learning_the_Face_Behind_a_Voice_CVPR_2019_paper.html

AUTHORS: Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T. Freeman, Michael Rubinstein, Wojciech Matusik

HIGHLIGHT: In this paper, we study the task of reconstructing a facial image of a person from a short audio recording of that person speaking.

TITLE: Joint Manifold Diffusion for Combining Predictions on Decoupled Observations

http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Joint_Manifold_Diffusion_for_Combining_Predictions_on_Decoupled_Observations_CVPR_2019_paper.html

AUTHORS: Kwang In Kim, Hyung Jin Chang

HIGHLIGHT: We present a new predictor combination algorithm that improves a given task predictor based on potentially relevant reference predictors.

TITLE: Audio Visual Scene-Aware Dialog

http://openaccess.thecvf.com/content_CVPR_2019/html/Alamri_Audio_Visual_Scene-Aware_Dialog_CVPR_2019_paper.html

AUTHORS: Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Anoop Cherian, Irfan Essa, Dhruv Batra, Tim K. Marks, Chiori Hori, Peter Anderson, Stefan Lee, Devi Parikh
HIGHLIGHT: We introduce the task of scene-aware dialog.

TITLE: Learning to Minify Photometric Stereo
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Learning_to_Minify_Photometric_Stereo_CVPR_2019_paper.html
AUTHORS: Junxuan Li, Antonio Robles-Kelly, Shaodi You, Yasuyuki Matsushita
HIGHLIGHT: We propose a method that can dramatically decrease the demands on the number of images by learning the most informative ones under different illumination conditions.

TITLE: Reflective and Fluorescent Separation Under Narrow-Band Illumination
http://openaccess.thecvf.com/content_CVPR_2019/html/Koyamatsu_Reflective_and_Fluorescent_Separation_Under_Narrow-Band_Illumination_CVPR_2019_paper.html
AUTHORS: Koji Koyamatsu, Daichi Hidaka, Takahiro Okabe, Hendrik P. A. Lensch
HIGHLIGHT: In this paper, we address the separation of reflective and fluorescent components in RGB images taken under narrow-band light sources such as LEDs.

TITLE: Depth From a Polarisation + RGB Stereo Pair
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhu_Depth_From_a_Polarisation_RGB_Stereo_Pair_CVPR_2019_paper.html
AUTHORS: Dizhong Zhu, William A. P. Smith
HIGHLIGHT: In this paper, we propose a hybrid depth imaging system in which a polarisation camera is augmented by a second image from a standard digital camera.

TITLE: Rethinking the Evaluation of Video Summaries
http://openaccess.thecvf.com/content_CVPR_2019/html/Otani_Rethinking_the_Evaluation_of_Video_Summaries_CVPR_2019_paper.html
AUTHORS: Mayu Otani, Yuta Nakashima, Esa Rahtu, Janne Heikkilä
HIGHLIGHT: In this paper, we will provide in-depth assessment of this pipeline using two popular benchmark datasets.

TITLE: What Object Should I Use? - Task Driven Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Sawatzky_What_Object_Should_I_Use_-_Task_Driven_Object_Detection_CVPR_2019_paper.html
AUTHORS: Johann Sawatzky, Yaser Souri, Christian Grund, Jurgen Gall
HIGHLIGHT: We therefore introduce the COCO-Tasks dataset which comprises about 40,000 images where the most suitable objects for 14 tasks have been annotated.

TITLE: Divergence Triangle for Joint Training of Generator Model, Energy-Based Model, and Inferential Model
http://openaccess.thecvf.com/content_CVPR_2019/html/Han_Divergence_Triangle_for_Joint_Training_of_Generator_Model_Energy-Based_Model_CVPR_2019_paper.html
AUTHORS: Tian Han, Erik Nijkamp, Xiaolin Fang, Mitch Hill, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT: This paper proposes the divergence triangle as a framework for joint training of a generator model, energy-based model and inference model.

TITLE: Image Deformation Meta-Networks for One-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Image_Deformation_Meta-Networks_for_One-Shot_Learning_CVPR_2019_paper.html
AUTHORS: Zitian Chen, Yanwei Fu, Yu-Xiong Wang, Lin Ma, Wei Liu, Martial Hebert
HIGHLIGHT: Our key insight is that, while the deformed images may not be visually realistic, they still maintain critical semantic information and contribute significantly to formulating classifier decision boundaries.

TITLE: Online High Rank Matrix Completion
http://openaccess.thecvf.com/content_CVPR_2019/html/Fan_Online_High_Rank_Matrix_Completion_CVPR_2019_paper.html
AUTHORS: Jicong Fan, Madeleine Udell
HIGHLIGHT: In this paper, we develop a new model for high rank matrix completion (HRMC), together with batch and online methods to fit the model and out-of-sample extension to complete new data.

TITLE: Multispectral Imaging for Fine-Grained Recognition of Powders on Complex Backgrounds
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhi_Multispectral_Imaging_for_Fine-Grained_Recognition_of_Powders_on_Complex_Backgrounds_CVPR_2019_paper.html
AUTHORS: Tiancheng Zhi, Bernardo R. Pires, Martial Hebert, Srinivasa G. Narasimhan

HIGHLIGHT: We present a method to select discriminative spectral bands to significantly reduce acquisition time while improving recognition accuracy.

TITLE: ContactDB: Analyzing and Predicting Grasp Contact via Thermal Imaging
http://openaccess.thecvf.com/content_CVPR_2019/html/Brahmbhatt_ContactDB_Analyzing_and_Predicting_Grasp_Contact_via_Thermal_Imaging_CVPR_2019_paper.html

AUTHORS: Samarth Brahmbhatt, Cusuh Ham, Charles C. Kemp, James Hays

HIGHLIGHT: We present ContactDB, a novel dataset of contact maps for household objects that captures the rich hand-object contact that occurs during grasping, enabled by use of a thermal camera.

TITLE: Robust Subspace Clustering With Independent and Piecewise Identically Distributed Noise Modeling
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Robust_Subspace_Clustering_With_Independent_and_Piecewise_Identically_Distributed_Noise_CVPR_2019_paper.html

AUTHORS: Yuanman Li, Jiantao Zhou, Xianwei Zheng, Jinyu Tian, Yuan Yan Tang

HIGHLIGHT: In this work, we propose an independent and piecewise identically distributed (i.p.i.d.) noise model, where the i.i.d. property only holds locally.

TITLE: What Correspondences Reveal About Unknown Camera and Motion Models?
http://openaccess.thecvf.com/content_CVPR_2019/html/Probst_What_Correspondences_Reveal_About_Unknown_Camera_and_Motion_Models_CVPR_2019_paper.html

AUTHORS: Thomas Probst, Ajad Chhatkuli, Danda Pani Paudel, Luc Van Gool

HIGHLIGHT: In this paper, we tackle this problem in two steps.

TITLE: Self-Calibrating Deep Photometric Stereo Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Self-Calibrating_Deep_Photometric_Stereo_Networks_CVPR_2019_paper.html

AUTHORS: Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, Kwan-Yee K. Wong

HIGHLIGHT: This paper proposes an uncalibrated photometric stereo method for non-Lambertian scenes based on deep learning.

TITLE: Argoverse: 3D Tracking and Forecasting With Rich Maps
http://openaccess.thecvf.com/content_CVPR_2019/html/Chang_Argoverse_3D_Tracking_and_Forecasting_With_Rich_Maps_CVPR_2019_paper.html

AUTHORS: Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, James Hays

HIGHLIGHT: We present Argoverse, a dataset designed to support autonomous vehicle perception tasks including 3D tracking and motion forecasting.

TITLE: Side Window Filtering
http://openaccess.thecvf.com/content_CVPR_2019/html/Yin_Side_Window_Filtering_CVPR_2019_paper.html

AUTHORS: Hui Yin, Yuanhao Gong, Guoping Qiu

HIGHLIGHT: Based on this insight, we propose a new Side Window Filtering (SWF) technique which aligns the window's side or corner with the pixel being processed.

TITLE: Defense Against Adversarial Images Using Web-Scale Nearest-Neighbor Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Dubey_Defense_Against_Adversarial_Images_Using_Web-Scale_Nearest-Neighbor_Search_CVPR_2019_paper.html

AUTHORS: Abhimanyu Dubey, Laurens van der Maaten, Zeki Yalniz, Yixuan Li, Dhruv Mahajan

HIGHLIGHT: In this work, we hypothesize that adversarial perturbations move the image away from the image manifold in the sense that there exists no physical process that could have produced the adversarial image.

TITLE: Incremental Object Learning From Contiguous Views
http://openaccess.thecvf.com/content_CVPR_2019/html/Stojanov_Incremental_Object_Learning_From_Contiguous_Views_CVPR_2019_paper.html

AUTHORS: Stefan Stojanov, Samarth Mishra, Ngoc Anh Thai, Nikhil Dhanda, Ahmad Humayun, Chen Yu, Linda B. Smith, James M. Rehg

HIGHLIGHT: In this work, we present CRIB (Continual Recognition Inspired by Babies), a synthetic incremental object learning environment that can produce data that models visual imagery produced by object exploration in early infancy.

TITLE: IP102: A Large-Scale Benchmark Dataset for Insect Pest Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_IP102_A_Large-Scale_Benchmark_Dataset_for_Insect_Pest_Recognition_CVPR_2019_paper.html
AUTHORS: Xiaoping Wu, Chi Zhan, Yu-Kun Lai, Ming-Ming Cheng, Jufeng Yang
HIGHLIGHT: In this paper, we collect a large-scale dataset named IP102 for insect pest recognition.

TITLE: CityFlow: A City-Scale Benchmark for Multi-Target Multi-Camera Vehicle Tracking and Re-Identification
http://openaccess.thecvf.com/content_CVPR_2019/html/Tang_CityFlow_A_City-Scale_Benchmark_for_Multi-Target_Multi-Camera_Vehicle_Tracking_and_CVPR_2019_paper.html
AUTHORS: Zheng Tang, Milind Naphade, Ming-Yu Liu, Xiaodong Yang, Stan Birchfield, Shuo Wang, Ratnesh Kumar, David Anastasiu, Jenq-Neng Hwang
HIGHLIGHT: This work introduces CityFlow, a city-scale traffic camera dataset consisting of more than 3 hours of synchronized HD videos from 40 cameras across 10 intersections, with the longest distance between two simultaneous cameras being 2.5 km.

TITLE: Social-IQ: A Question Answering Benchmark for Artificial Social Intelligence
http://openaccess.thecvf.com/content_CVPR_2019/html/Zadeh_Social-IQ_A_Question_Answering_Benchmark_for_Artificial_Social_Intelligence_CVPR_2019_paper.html
AUTHORS: Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, Louis-Philippe Morency
HIGHLIGHT: In this paper, we introduce Social-IQ, a unconstrained benchmark specifically designed to train and evaluate socially intelligent technologies.

TITLE: On Zero-Shot Recognition of Generic Objects
http://openaccess.thecvf.com/content_CVPR_2019/html/Hascoet_On_Zero-Shot_Recognition_of_Generic_Objects_CVPR_2019_paper.html
AUTHORS: Tristan Hascoet, Yasuo Ariki, Tetsuya Takiguchi
HIGHLIGHT: In this paper, we argue that the main reason behind this apparent lack of progress is the poor quality of this benchmark.

TITLE: Explicit Bias Discovery in Visual Question Answering Models
http://openaccess.thecvf.com/content_CVPR_2019/html/Manjunatha_Explicit_Bias_Discovery_in_Visual_Question_Answering_Models_CVPR_2019_paper.html
AUTHORS: Varun Manjunatha, Nirat Saini, Larry S. Davis
HIGHLIGHT: Researchers have observed that Visual Question Answering (VQA) models tend to answer questions by learning statistical biases in the data.

TITLE: REPAIR: Removing Representation Bias by Dataset Resampling
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_REPAIR_Removing_Representation_Bias_by_Dataset_Resampling_CVPR_2019_paper.html
AUTHORS: Yi Li, Nuno Vasconcelos
HIGHLIGHT: REPAIR: Removing Representation Bias by Dataset Resampling

TITLE: Label Efficient Semi-Supervised Learning via Graph Filtering
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Label_Efficient_Semi-Supervised_Learning_via_Graph_Filtering_CVPR_2019_paper.html
AUTHORS: Qimai Li, Xiao-Ming Wu, Han Liu, Xiaotong Zhang, Zhichao Guan
HIGHLIGHT: In this paper, we address label efficient semi-supervised learning from a graph filtering perspective.

TITLE: MVTec AD -- A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Bergmann_MVTec_AD--A_Comprehensive_Real-World_Dataset_for_Unsupervised_Anomaly_CVPR_2019_paper.html
AUTHORS: Paul Bergmann, Michael Fauser, David Sattlegger, Carsten Steger
HIGHLIGHT: We introduce the MVTec Anomaly Detection (MVTec AD) dataset containing 5354 high-resolution color images of different object and texture categories.

TITLE: ABC: A Big CAD Model Dataset for Geometric Deep Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Koch_ABC_A_Big_CAD_Model_Dataset_for_Geometric_Deep_Learning_CVPR_2019_paper.html
AUTHORS: Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, Daniele Panozzo
HIGHLIGHT: We introduce ABC-Dataset, a collection of one million Computer-Aided Design (CAD) models for research of geometric deep learning methods and applications.

TITLE: Tightness-Aware Evaluation Protocol for Scene Text Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Tightness-Aware_Evaluation_Protocol_for_Scene_Text_Detection_CVPR_2019_paper.html
AUTHORS: Yuliang Liu, Lianwen Jin, Zecheng Xie, Canjie Luo, Shuaitao Zhang, Lele Xie
HIGHLIGHT: Therefore, this paper proposes a novel evaluation protocol called Tightness-aware Intersect-over-Union (TIOU) metric that could quantify completeness of ground truth, compactness of detection, and tightness of matching degree.

TITLE: Local Features and Visual Words Emerge in Activations
http://openaccess.thecvf.com/content_CVPR_2019/html/Simeoni_Local_Features_and_Visual_Words_Emerge_in_Activations_CVPR_2019_paper.html
AUTHORS: Oriane Simeoni, Yannis Avrithis, Ondrej Chum
HIGHLIGHT: We propose a novel method of deep spatial matching (DSM) for image retrieval.

TITLE: Hyperspectral Image Super-Resolution With Optimized RGB Guidance
http://openaccess.thecvf.com/content_CVPR_2019/html/Fu_Hyperspectral_Image_Super-Resolution_With_Optimized_RGB_Guidance_CVPR_2019_paper.html
AUTHORS: Ying Fu, Tao Zhang, Yinqiang Zheng, Debing Zhang, Hua Huang
HIGHLIGHT: Previous methods for this fusion task usually employ hand-crafted priors to model the underlying structure of the latent high resolution HSI, and the effect of the camera spectral response (CSR) of the RGB camera on super-resolution accuracy has rarely been investigated.

TITLE: Adaptive Confidence Smoothing for Generalized Zero-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Atzmon_Adaptive_Confidence_Smoothing_for_Generalized_Zero-Shot_Learning_CVPR_2019_paper.html
AUTHORS: Yuval Atzmon, Gal Chechik
HIGHLIGHT: Here we describe a probabilistic approach that breaks the model into three modular components, and then combines them in a consistent way.

TITLE: PMS-Net: Robust Haze Removal Based on Patch Map for Single Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_PMS-Net_Robust_Haze_Removal_Based_on_Patch_Map_for_Single_CVPR_2019_paper.html
AUTHORS: Wei-Ting Chen, Jian-Jiun Ding, Sy-Yen Kuo
HIGHLIGHT: In this paper, we proposed a novel haze removal algorithm based on a new feature called the patch map.

TITLE: Deep Spherical Quantization for Image Search
http://openaccess.thecvf.com/content_CVPR_2019/html/Eghbali_Deep_Spherical_Quantization_for_Image_Search_CVPR_2019_paper.html
AUTHORS: Sepehr Eghbali, Ladan Tahvildari
HIGHLIGHT: In this paper, we put forward Deep Spherical Quantization (DSQ), a novel method to make deep convolutional neural networks generate supervised and compact binary codes for efficient image search.

TITLE: Large-Scale Interactive Object Segmentation With Human Annotators
http://openaccess.thecvf.com/content_CVPR_2019/html/Benenson_Large-Scale_Interactive_Object_Segmentation_With_Human_Annotators_CVPR_2019_paper.html
AUTHORS: Rodrigo Benenson, Stefan Popov, Vittorio Ferrari
HIGHLIGHT: In this paper we make several contributions to interactive segmentation: (1) we systematically explore in simulation the design space of deep interactive segmentation models and report new insights and caveats; (2) we execute a large-scale annotation campaign with real human annotators, producing masks for 2.5M instances on the OpenImages dataset.

TITLE: A Poisson-Gaussian Denoising Dataset With Real Fluorescence Microscopy Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_A_Poisson-Gaussian_Denoising_Dataset_With_Real_Fluorescence_Microscopy_Images_CVPR_2019_paper.html
AUTHORS: Yide Zhang, Yin hao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, Scott Howard
HIGHLIGHT: In this paper, we fill this gap by constructing a dataset - the Fluorescence Microscopy Denoising (FMD) dataset - that is dedicated to Poisson-Gaussian denoising.

TITLE: Task Agnostic Meta-Learning for Few-Shot Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Jamal_Task_Agnostic_Meta-Learning_for_Few-Shot_Learning_CVPR_2019_paper.html
AUTHORS: Muhammad Abdullah Jamal, Guo-Jun Qi

HIGHLIGHT: Specifically, we present an entropy-based approach that meta-learns an unbiased initial model with the largest uncertainty over the output labels by preventing it from over-performing in classification tasks.

TITLE: Progressive Ensemble Networks for Zero-Shot Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Ye_Progressive_Ensemble_Networks_for_Zero-Shot_Recognition_CVPR_2019_paper.html

AUTHORS: Meng Ye, Yuhong Guo

HIGHLIGHT: In this paper, we propose a novel progressive ensemble network model with multiple projected label embeddings to address zero-shot image recognition.

TITLE: Direct Object Recognition Without Line-Of-Sight Using Optical Coherence

http://openaccess.thecvf.com/content_CVPR_2019/html/Lei_Direct_Object_Recognition_Without_Line-Of-Sight_Using_Optical_Coherence_CVPR_2019_paper.html

AUTHORS: Xin Lei, Liangyu He, Yixuan Tan, Ken Xingze Wang, Xinggang Wang, Yihan Du, Shanhui Fan, Zongfu Yu

HIGHLIGHT: We introduce a novel approach based on speckle pattern recognition with deep neural network, which is simpler and more robust than other NLOS recognition methods.

TITLE: Atlas of Digital Pathology: A Generalized Hierarchical Histological Tissue Type-Annotated Database for Deep Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Hosseini_Atlas_of_Digital_Pathology_A_Generalized_Hierarchical_Histological_Tissue_Type-Annotated_CVPR_2019_paper.html

AUTHORS: Mahdi S. Hosseini, Lyndon Chan, Gabriel Tse, Michael Tang, Jun Deng, Sajad Norouzi, Corwyn Rowsell, Konstantinos N. Plataniotis, Savvas Damaskinos

HIGHLIGHT: In this paper, we propose a new digital pathology database, the "Atlas of Digital Pathology" (or ADP), which comprises of 17,668 patch images extracted from 100 slides annotated with up to 57 hierarchical HTTs.