

TITLE: SDRSAC: Semidefinite-Based Randomized Approach for Robust Point Cloud Registration Without Correspondences
http://openaccess.thecvf.com/content_CVPR_2019/html/Le_SDRSAC_Semidefinite-Based_Randomized_Approach_for_Robust_Point_Cloud_Registration_Without_CVPR_2019_paper.html
AUTHORS: Huu M. Le, Thanh-Toan Do, Tuan Hoang, Ngai-Man Cheung
HIGHLIGHT: This paper presents a novel randomized algorithm for robust point cloud registration without correspondences.

TITLE: BAD SLAM: Bundle Adjusted Direct RGB-D SLAM
http://openaccess.thecvf.com/content_CVPR_2019/html/Schops_BAD_SLAM_Bundle_Adjusted_Direct_RGB-D_SLAM_CVPR_2019_paper.html
AUTHORS: Thomas Schops, Torsten Sattler, Marc Pollefeys
HIGHLIGHT: In contrast, in this paper we present a novel, fast direct BA formulation which we implement in a real-time dense RGB-D SLAM algorithm.

TITLE: Revealing Scenes by Inverting Structure From Motion Reconstructions
http://openaccess.thecvf.com/content_CVPR_2019/html/Pittaluga_Revealing_Scenes_by_Inverting_Structure_From_Motion_Reconstructions_CVPR_2019_paper.html
AUTHORS: Francesco Pittaluga, Sanjeev J. Koppal, Sing Bing Kang, Sudipta N. Sinha
HIGHLIGHT: In this paper, we show, for the first time, that such point clouds retain enough information to reveal scene appearance and compromise privacy.

TITLE: Strand-Accurate Multi-View Hair Capture
http://openaccess.thecvf.com/content_CVPR_2019/html/Nam_Strand-Accurate_Multi-View_Hair_Capture_CVPR_2019_paper.html
AUTHORS: Giljoo Nam, Chenglei Wu, Min H. Kim, Yaser Sheikh
HIGHLIGHT: In this paper, we present the first method to capture high-fidelity hair geometry with strand-level accuracy.

TITLE: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation
http://openaccess.thecvf.com/content_CVPR_2019/html/Park_DeepSDF_Learning_Continuous_Signed_Distance_Functions_for_Shape_Representation_CVPR_2019_paper.html
AUTHORS: Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, Steven Lovegrove
HIGHLIGHT: In this work, we introduce DeepSDF, a learned continuous Signed Distance Function (SDF) representation of a class of shapes that enables high quality shape representation, interpolation and completion from partial and noisy 3D input data.

TITLE: Pushing the Boundaries of View Extrapolation With Multiplane Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Srinivasan_Pushing_the_Boundaries_of_View_Extrapolation_With_Multiplane_Images_CVPR_2019_paper.html
AUTHORS: Pratul P. Srinivasan, Richard Tucker, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng, Noah Snavely
HIGHLIGHT: We explore the problem of view synthesis from a narrow baseline pair of images, and focus on generating high-quality view extrapolations with plausible disocclusions.

TITLE: GA-Net: Guided Aggregation Net for End-To-End Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_GA-Net_Guided_Aggregation_Net_for_End-To-End_Stereo_Matching_CVPR_2019_paper.html
AUTHORS: Feihu Zhang, Victor Prisacariu, Ruigang Yang, Philip H.S. Torr
HIGHLIGHT: We propose two novel neural net layers, aimed at capturing local and the whole-image cost dependencies respectively.

TITLE: Real-Time Self-Adaptive Deep Stereo
http://openaccess.thecvf.com/content_CVPR_2019/html/Tonioni_Real-Time_Self-Adaptive_Deep_Stereo_CVPR_2019_paper.html
AUTHORS: Alessio Tonioni, Fabio Tosi, Matteo Poggi, Stefano Mattoccia, Luigi Di Stefano
HIGHLIGHT: Instead, we propose to perform unsupervised and continuous online adaptation of a deep stereo network, which allows for preserving its accuracy in any environment.

TITLE: LAF-Net: Locally Adaptive Fusion Networks for Stereo Confidence Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_LAF-Net_Locally_Adaptive_Fusion_Networks_for_Stereo_Confidence_Estimation_CVPR_2019_paper.html
AUTHORS: Sunok Kim, Seungryong Kim, Dongbo Min, Kwanghoon Sohn
HIGHLIGHT: We present a novel method that estimates confidence map of an initial disparity by making full use of tri-modal input, including matching cost, disparity, and color image through deep networks.

TITLE: NM-Net: Mining Reliable Neighbors for Robust Feature Correspondences

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_NM-Net_Mining_Reliable_Neighbors_for_Robust_Feature_Correspondences_CVPR_2019_paper.html
AUTHORS: Chen Zhao, Zhiguo Cao, Chi Li, Xin Li, Jiaqi Yang
HIGHLIGHT: To address this issue, we present a compatibility-specific mining method to search for consistent neighbors.

TITLE: Coordinate-Free Carlsson-Weinshall Duality and Relative Multi-View Geometry
http://openaccess.thecvf.com/content_CVPR_2019/html/Trager_Coordinate-Free_Carlsson-Weinshall_Duality_and_Relative_Multi-View_Geometry_CVPR_2019_paper.html
AUTHORS: Matthew Trager, Martial Hebert, Jean Ponce
HIGHLIGHT: We present a coordinate-free description of Carlsson-Weinshall duality between scene points and camera pinholes and use it to derive a new characterization of primal/dual multi-view geometry.

TITLE: Deep Reinforcement Learning of Volume-Guided Progressive View Inpainting for 3D Point Scene Completion From a Single Depth Image
http://openaccess.thecvf.com/content_CVPR_2019/html/Han_Deep_Reinforcement_Learning_of_Volume-Guided_Progressive_View_Inpainting_for_3D_CVPR_2019_paper.html
AUTHORS: Xiaoguang Han, Zhaoxuan Zhang, Dong Du, Mingdai Yang, Jingming Yu, Pan Pan, Xin Yang, Ligang Liu, Zixiang Xiong, Shuguang Cui
HIGHLIGHT: We present a deep reinforcement learning method of progressive view inpainting for 3D point scene completion under volume guidance, achieving high-quality scene reconstruction from only a single depth image with severe occlusion.

TITLE: Structural Relational Reasoning of Point Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Duan_Structural_Relational_Reasoning_of_Point_Clouds_CVPR_2019_paper.html
AUTHORS: Yueqi Duan, Yu Zheng, Jiwen Lu, Jie Zhou, Qi Tian
HIGHLIGHT: In this paper, we propose an effective plug-and-play module called the structural relation network (SRN) to reason about the structural dependencies of local regions in 3D point clouds.

TITLE: MVF-Net: Multi-View 3D Face Morphable Model Regression
http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_MVF-Net_Multi-View_3D_Face_Morphable_Model_Regression_CVPR_2019_paper.html
AUTHORS: Fanzi Wu, Linchao Bao, Yajing Chen, Yonggen Ling, Yibing Song, Songnan Li, King Ng Ngan, Wei Liu
HIGHLIGHT: We in this paper explore 3DMM-based shape recovery in a different setting, where a set of multi-view facial images are given as input.

TITLE: Photometric Mesh Optimization for Video-Aligned 3D Object Reconstruction
http://openaccess.thecvf.com/content_CVPR_2019/html/Lin_Photometric_Mesh_Optimization_for_Video-Aligned_3D_Object_Reconstruction_CVPR_2019_paper.html
AUTHORS: Chen-Hsuan Lin, Oliver Wang, Bryan C. Russell, Eli Shechtman, Vladimir G. Kim, Matthew Fisher, Simon Lucey
HIGHLIGHT: In this paper, we address the problem of 3D object mesh reconstruction from RGB videos.

TITLE: Guided Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2019/html/Poggi_Guided_Stereo_Matching_CVPR_2019_paper.html
AUTHORS: Matteo Poggi, Davide Pallotti, Fabio Tosi, Stefano Mattoccia
HIGHLIGHT: Therefore, in this paper, we introduce Guided Stereo Matching, a novel paradigm leveraging a small amount of sparse, yet reliable depth measurements retrieved from an external source enabling to ameliorate this weakness.

TITLE: Unsupervised Event-Based Learning of Optical Flow, Depth, and Egomotion
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhu_Unsupervised_Event-Based_Learning_of_Optical_Flow_Depth_and_Egomotion_CVPR_2019_paper.html
AUTHORS: Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, Kostas Daniilidis
HIGHLIGHT: In this work, we propose a novel framework for unsupervised learning for event cameras that learns motion information from only the event stream.

TITLE: Modeling Local Geometric Structure of 3D Point Clouds Using Geo-CNN
http://openaccess.thecvf.com/content_CVPR_2019/html/Lan_Modeling_Local_Geometric_Structure_of_3D_Point_Clouds_Using_Geo-CNN_CVPR_2019_paper.html
AUTHORS: Shiyi Lan, Ruichi Yu, Gang Yu, Larry S. Davis
HIGHLIGHT: To address this problem, we propose Geo-CNN, which applies a generic convolution-like operation dubbed as GeoConv to each point and its local neighborhood.

TITLE: 3D Local Features for Direct Pairwise Registration
http://openaccess.thecvf.com/content_CVPR_2019/html/Deng_3D_Local_Features_for_Direct_Pairwise_Registration_CVPR_2019_paper.html
AUTHORS: Haowen Deng, Tolga Birdal, Slobodan Ilic
HIGHLIGHT: We present a novel, data driven approach for solving the problem of registration of two point cloud scans.

TITLE: HPLFlowNet: Hierarchical Permutohedral Lattice FlowNet for Scene Flow Estimation on Large-Scale Point Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Gu_HPLFlowNet_Hierarchical_Permutohedral_Lattice_FlowNet_for_Scene_Flow_Estimation_on_CVPR_2019_paper.html
AUTHORS: Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, Panqu Wang
HIGHLIGHT: We present a novel deep neural network architecture for end-to-end scene flow estimation that directly operates on large-scale 3D point clouds.

TITLE: GPSfM: Global Projective SFM Using Algebraic Constraints on Multi-View Fundamental Matrices
http://openaccess.thecvf.com/content_CVPR_2019/html/Kasten_GPSfM_Global_Projective_SFM_Using_Algebraic_Constraints_on_Multi-View_Fundamental_CVPR_2019_paper.html
AUTHORS: Yoni Kasten, Amnon Geifman, Meirav Galun, Ronen Basri
HIGHLIGHT: This paper addresses the problem of recovering projective camera matrices from collections of fundamental matrices in multiview settings.

TITLE: Group-Wise Correlation Stereo Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Guo_Group-Wise_Correlation_Stereo_Network_CVPR_2019_paper.html
AUTHORS: Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, Hongsheng Li
HIGHLIGHT: In this paper, we propose to construct the cost volume by group-wise correlation.

TITLE: Multi-Level Context Ultra-Aggregation for Stereo Matching
http://openaccess.thecvf.com/content_CVPR_2019/html/Nie_Multi-Level_Context_Ultra-Aggregation_for_Stereo_Matching_CVPR_2019_paper.html
AUTHORS: Guang-Yu Nie, Ming-Ming Cheng, Yun Liu, Zhengfa Liang, Deng-Ping Fan, Yue Liu, Yongtian Wang
HIGHLIGHT: In this paper, we propose a unary features descriptor using multi-level context ultra-aggregation (MCUA), which encapsulates all convolutional features into a more discriminative representation by intra- and inter-level features combination.

TITLE: Large-Scale, Metric Structure From Motion for Unordered Light Fields
http://openaccess.thecvf.com/content_CVPR_2019/html/Nousias_Large-Scale_Metric_Structure_From_Motion_for_Unordered_Light_Fields_CVPR_2019_paper.html
AUTHORS: Sotiris Nousias, Manolis Lourakis, Christos Bergeles
HIGHLIGHT: This paper presents a large scale, metric Structure from Motion (SfM) pipeline for generalised cameras with overlapping fields-of-view, and demonstrates it using Light Field (LF) images.

TITLE: Understanding the Limitations of CNN-Based Absolute Camera Pose Regression
http://openaccess.thecvf.com/content_CVPR_2019/html/Sattler_Understanding_the_Limitations_of_CNN-Based_Absolute_Camera_Pose_Regression_CVPR_2019_paper.html
AUTHORS: Torsten Sattler, Qunjie Zhou, Marc Pollefeys, Laura Leal-Taixe
HIGHLIGHT: To understand this behavior, we develop a theoretical model for camera pose regression.

TITLE: DeepLiDAR: Deep Surface Normal Guided Depth Prediction for Outdoor Scene From Sparse LiDAR Data and Single Color Image
http://openaccess.thecvf.com/content_CVPR_2019/html/Qiu_DeepLiDAR_Deep_Surface_Normal_Guided_Depth_Prediction_for_Outdoor_Scene_CVPR_2019_paper.html
AUTHORS: Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, Marc Pollefeys
HIGHLIGHT: In this paper, we propose a deep learning architecture that produces accurate dense depth for the outdoor scene from a single color image and a sparse depth.

TITLE: Modeling Point Clouds With Self-Attention and Gumbel Subset Sampling
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Modeling_Point_Clouds_With_Self-Attention_and_Gumbel_Subset_Sampling_CVPR_2019_paper.html
AUTHORS: Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, Qi Tian
HIGHLIGHT: Thereby, we for the first time propose an end-to-end learnable and task-agnostic sampling operation, named Gumbel Subset Sampling (GSS), to select a representative subset of input points.

TITLE: Learning With Batch-Wise Optimal Transport Loss for 3D Shape Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_Learning_With_Batch-Wise_Optimal_Transport_Loss_for_3D_Shape_Recognition_CVPR_2019_paper.html
AUTHORS: Lin Xu, Han Sun, Yuai Liu
HIGHLIGHT: In this paper, we show how to learn an importance-driven distance metric via optimal transport programming from batches of samples.

TITLE: DenseFusion: 6D Object Pose Estimation by Iterative Dense Fusion
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_DenseFusion_6D_Object_Pose_Estimation_by_Iterative_Dense_Fusion_CVPR_2019_paper.html
AUTHORS: Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martin-Martin, Cewu Lu, Li Fei-Fei, Silvio Savarese
HIGHLIGHT: In this work, we present DenseFusion, a generic framework for estimating 6D pose of a set of known objects from RGB-D images.

TITLE: Multi-Scale Geometric Consistency Guided Multi-View Stereo
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_Multi-Scale_Geometric_Consistency_Guided_Multi-View_Stereo_CVPR_2019_paper.html
AUTHORS: Qingshan Xu, Wenbing Tao
HIGHLIGHT: In this paper, we propose an efficient multi-scale geometric consistency guided multi-view stereo method for accurate and complete depth map estimation.

TITLE: Privacy Preserving Image-Based Localization
http://openaccess.thecvf.com/content_CVPR_2019/html/Speciale_Privacy_Preserving_Image-Based_Localization_CVPR_2019_paper.html
AUTHORS: Pablo Speciale, Johannes L. Schonberger, Sing Bing Kang, Sudipta N. Sinha, Marc Pollefeys
HIGHLIGHT: This paper proposes the first solution to what we call privacy preserving image-based localization.

TITLE: SimulCap : Single-View Human Performance Capture With Cloth Simulation
http://openaccess.thecvf.com/content_CVPR_2019/html/Yu_SimulCap_Single-View_Human_Performance_Capture_With_Cloth_Simulation_CVPR_2019_paper.html
AUTHORS: Tao Yu, Zerong Zheng, Yuan Zhong, Jianhui Zhao, Qionghai Dai, Gerard Pons-Moll, Yebin Liu
HIGHLIGHT: This paper proposes a new method for live free-viewpoint human performance capture with dynamic details (e.g., cloth wrinkles) using a single RGBD camera.

TITLE: Hierarchical Deep Stereo Matching on High-Resolution Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Hierarchical_Deep_Stereo_Matching_on_High-Resolution_Images_CVPR_2019_paper.html
AUTHORS: Gengshan Yang, Joshua Manela, Michael Happold, Deva Ramanan
HIGHLIGHT: To address this issue, we propose an end-to-end framework that searches for correspondences incrementally over a coarse-to-fine hierarchy.

TITLE: Recurrent MVSNet for High-Resolution Multi-View Stereo Depth Inference
http://openaccess.thecvf.com/content_CVPR_2019/html/Yao_Recurrent_MVSNet_for_High-Resolution_Multi-View_Stereo_Depth_Inference_CVPR_2019_paper.html
AUTHORS: Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, Long Quan
HIGHLIGHT: In this paper, we introduce a scalable multi-view stereo framework based on the recurrent neural network.

TITLE: Synthesizing 3D Shapes From Silhouette Image Collections Using Multi-Projection Generative Adversarial Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Synthesizing_3D_Shapes_From_Silhouette_Image_Collections_Using_Multi-Projection_Generative_CVPR_2019_paper.html
AUTHORS: Xiao Li, Yue Dong, Pieter Peers, Xin Tong
HIGHLIGHT: We present a new weakly supervised learning-based method for generating novel category-specific 3D shapes from unoccluded image collections.

TITLE: The Perfect Match: 3D Point Cloud Matching With Smoothed Densities
http://openaccess.thecvf.com/content_CVPR_2019/html/Gojcic_The_Perfect_Match_3D_Point_Cloud_Matching_With_Smoothed_Densities_CVPR_2019_paper.html
AUTHORS: Zan Gojcic, Caifa Zhou, Jan D. Wegner, Andreas Wieser
HIGHLIGHT: We propose 3DSmoothNet, a full workflow to match 3D point clouds with a siamese deep learning architecture and fully convolutional layers using a voxelized smoothed density value (SDV) representation.

TITLE: Recurrent Neural Network for (Un-)Supervised Learning of Monocular Video Visual Odometry and Depth
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Recurrent_Neural_Network_for_Un-Supervised_Learning_of_Monocular_Video_Visual_CVPR_2019_paper.html
AUTHORS: Rui Wang, Stephen M. Pizer, Jan-Michael Frahm
HIGHLIGHT: We propose a learning-based, multi-view dense depth map and odometry estimation method that uses Recurrent Neural Networks (RNN) and trains utilizing multi-view image reprojection and forward-backward flow-consistency losses.

TITLE: PointWeb: Enhancing Local Neighborhood Features for Point Cloud Processing
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_PointWeb_Enhancing_Local_Neighborhood_Features_for_Point_Cloud_Processing_CVPR_2019_paper.html
AUTHORS: Hengshuang Zhao, Li Jiang, Chi-Wing Fu, Jiaya Jia
HIGHLIGHT: This paper presents PointWeb, a new approach to extract contextual features from local neighborhood in a point cloud.

TITLE: Scan2Mesh: From Unstructured Range Scans to 3D Meshes
http://openaccess.thecvf.com/content_CVPR_2019/html/Dai_Scan2Mesh_From_Unstructured_Range_Scans_to_3D_Meshes_CVPR_2019_paper.html
AUTHORS: Angela Dai, Matthias Niessner
HIGHLIGHT: We introduce Scan2Mesh, a novel data-driven generative approach which transforms an unstructured and potentially incomplete range scan into a structured 3D mesh representation.

TITLE: Unsupervised Domain Adaptation for ToF Data Denoising With Adversarial Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Agresti_Unsupervised_Domain_Adaptation_for_ToF_Data_Denoising_With_Adversarial_Learning_CVPR_2019_paper.html
AUTHORS: Gianluca Agresti, Henrik Schaefer, Piergiorgio Sartor, Pietro Zanuttigh
HIGHLIGHT: In this paper, we avoid to rely on labeled real data in the learning framework.

TITLE: Learning Independent Object Motion From Unlabelled Stereoscopic Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Cao_Learning_Independent_Object_Motion_From_Unlabelled_Stereoscopic_Videos_CVPR_2019_paper.html
AUTHORS: Zhe Cao, Abhishek Kar, Christian Hane, Jitendra Malik
HIGHLIGHT: We present a system for learning motion maps of independently moving objects from stereo videos.

TITLE: Triangulation Learning Network: From Monocular to Stereo 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Qin_Triangulation_Learning_Network_From_Monocular_to_Stereo_3D_Object_Detection_CVPR_2019_paper.html
AUTHORS: Zengyi Qin, Jinglu Wang, Yan Lu
HIGHLIGHT: In this paper, we study the problem of 3D object detection from stereo images, in which the key challenge is how to effectively utilize stereo information.

TITLE: Connecting the Dots: Learning Representations for Active Monocular Depth Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Riegler_Connecting_the_Dots_Learning_Representations_for_Active_Monocular_Depth_Estimation_CVPR_2019_paper.html
AUTHORS: Gernot Riegler, Yiyi Liao, Simon Donne, Vladlen Koltun, Andreas Geiger
HIGHLIGHT: We propose a technique for depth estimation with a monocular structured-light camera, i.e., a calibrated stereo set-up with one camera and one laser projector.

TITLE: Learning Non-Volumetric Depth Fusion Using Successive Reprojections
http://openaccess.thecvf.com/content_CVPR_2019/html/Donne_Learning_Non-Volumetric_Depth_Fusion_Using_Successive_Reprojections_CVPR_2019_paper.html
AUTHORS: Simon Donne, Andreas Geiger
HIGHLIGHT: In this work we propose to learn an auto-regressive depth refinement directly from data.

TITLE: Stereo R-CNN Based 3D Object Detection for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Stereo_R-CNN_Based_3D_Object_Detection_for_Autonomous_Driving_CVPR_2019_paper.html
AUTHORS: Peiliang Li, Xiaozhi Chen, Shaojie Shen
HIGHLIGHT: We propose a 3D object detection method for autonomous driving by fully exploiting the sparse and dense, semantic and geometry information in stereo imagery.

TITLE: Hybrid Scene Compression for Visual Localization

http://openaccess.thecvf.com/content_CVPR_2019/html/Camposeco_Hybrid_Scene_Compression_for_Visual_Localization_CVPR_2019_paper.html

AUTHORS: Federico Camposeco, Andrea Cohen, Marc Pollefeys, Torsten Sattler
HIGHLIGHT: In this work, we introduce a new hybrid compression algorithm that uses a given memory limit in a more effective way.

TITLE: PointConv: Deep Convolutional Networks on 3D Point Clouds

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_PointConv_Deep_Convolutional_Networks_on_3D_Point_Clouds_CVPR_2019_paper.html

AUTHORS: Wenxuan Wu, Zhongang Qi, Li Fuxin
HIGHLIGHT: In this paper, we extend the dynamic filter to a new convolution operation, named PointConv.

TITLE: Octree Guided CNN With Spherical Kernels for 3D Point Clouds

http://openaccess.thecvf.com/content_CVPR_2019/html/Lei_Octree_Guided_CNN_With_Spherical_Kernels_for_3D_Point_Clouds_CVPR_2019_paper.html

AUTHORS: Huan Lei, Naveed Akhtar, Ajmal Mian
HIGHLIGHT: We propose an octree guided neural network architecture and spherical convolutional kernel for machine learning from arbitrary 3D point clouds.

TITLE: VITAMIN-E: Visual Tracking and Mapping With Extremely Dense Feature Points

http://openaccess.thecvf.com/content_CVPR_2019/html/Yokozuka_VITAMIN-E_Visual_Tracking_and_Mapping_With_Extremely_Dense_Feature_Points_CVPR_2019_paper.html

AUTHORS: Masashi Yokozuka, Shuji Oishi, Simon Thompson, Atsuhiko Banno
HIGHLIGHT: In this paper, we propose a novel indirect monocular simultaneous localization and mapping (SLAM) algorithm called "VITAMIN-E," which is highly accurate and robust as a result of tracking extremely dense feature points.

TITLE: Conditional Single-View Shape Generation for Multi-View Stereo Reconstruction

http://openaccess.thecvf.com/content_CVPR_2019/html/Wei_Conditional_Single-View_Shape_Generation_for_Multi-View_Stereo_Reconstruction_CVPR_2019_paper.html

AUTHORS: Yi Wei, Shaohui Liu, Wang Zhao, Jiwen Lu
HIGHLIGHT: In this paper, we present a new perspective towards image-based shape generation.

TITLE: Learning to Adapt for Stereo

http://openaccess.thecvf.com/content_CVPR_2019/html/Tonioni_Learning_to_Adapt_for_Stereo_CVPR_2019_paper.html

AUTHORS: Alessio Tonioni, Oscar Rahnama, Thomas Joy, Luigi Di Stefano, Thalaiyasingam Ajanthan, Philip H.S. Torr
HIGHLIGHT: In this work, we introduce a "learning-to-adapt" framework that enables deep stereo methods to continuously adapt to new target domains in an unsupervised manner.

TITLE: 3D Appearance Super-Resolution With Deep Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_3D_Appearance_Super-Resolution_With_Deep_Learning_CVPR_2019_paper.html

AUTHORS: Yawei Li, Vagia Tsiminaki, Radu Timofte, Marc Pollefeys, Luc Van Gool
HIGHLIGHT: We introduce a 3D appearance SR (3DASR) dataset based on the existing ETH3D [42], SyB3R [31], MiddleBury, and our Collection of 3D scenes from TUM [21], Fountain [51] and Relief [53].

TITLE: Radial Distortion Triangulation

http://openaccess.thecvf.com/content_CVPR_2019/html/Kukelova_Radial_Distortion_Triangulation_CVPR_2019_paper.html

AUTHORS: Zuzana Kukelova, Viktor Larsson
HIGHLIGHT: This paper presents the first optimal, maximal likelihood, solution to the triangulation problem for radially distorted cameras.

TITLE: Robust Point Cloud Based Reconstruction of Large-Scale Outdoor Scenes

http://openaccess.thecvf.com/content_CVPR_2019/html/Lan_Robust_Point_Cloud_Based_Reconstruction_of_Large-Scale_Outdoor_Scenes_CVPR_2019_paper.html

AUTHORS: Ziquan Lan, Zi Jian Yew, Gim Hee Lee
HIGHLIGHT: To alleviate this problem, we propose a probabilistic approach for robust back-end optimization in the presence of outliers.

TITLE: Perturbation Analysis of the 8-Point Algorithm: A Case Study for Wide FoV Cameras

http://openaccess.thecvf.com/content_CVPR_2019/html/da_Silveira_Perturbation_Analysis_of_the_8-Point_Algorithm_A_Case_Study_for_CVPR_2019_paper.html

AUTHORS: Thiago L. T. da Silveira, Claudio R. Jung
HIGHLIGHT: This paper presents a perturbation analysis for the estimate of epipolar matrices using the 8-Point Algorithm (8-PA).

TITLE: Robustness of 3D Deep Learning in an Adversarial Setting
http://openaccess.thecvf.com/content_CVPR_2019/html/Wicker_Robustness_of_3D_Deep_Learning_in_an_Adversarial_Setting_CVPR_2019_paper.html

AUTHORS: Matthew Wicker, Marta Kwiatkowska
HIGHLIGHT: In this work, we develop an algorithm for analysis of pointwise robustness of neural networks that operate on 3D data.

TITLE: SceneCode: Monocular Dense Semantic Reconstruction Using Learned Encoded Scene Representations
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhi_SceneCode_Monocular_Dense_Semantic_Reconstruction_Using_Learned_Encoded_Scene_Representations_CVPR_2019_paper.html

AUTHORS: Shuaifeng Zhi, Michael Bloesch, Stefan Leutenegger, Andrew J. Davison
HIGHLIGHT: We introduce a new compact and optimisable semantic representation by training a variational auto-encoder that is conditioned on a colour image.

TITLE: StereoDRNet: Dilated Residual StereoNet
http://openaccess.thecvf.com/content_CVPR_2019/html/Chabra_StereoDRNet_Dilated_Residual_StereoNet_CVPR_2019_paper.html

AUTHORS: Rohan Chabra, Julian Straub, Christopher Sweeney, Richard Newcombe, Henry Fuchs
HIGHLIGHT: We propose a system that uses a convolution neural network (CNN) to estimate depth from a stereo pair followed by volumetric fusion of the predicted depth maps to produce a 3D reconstruction of a scene.

TITLE: The Alignment of the Spheres: Globally-Optimal Spherical Mixture Alignment for Camera Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Campbell_The_Alignment_of_the_Spheres_Globally-Optimal_Spherical_Mixture_Alignment_for_CVPR_2019_paper.html

AUTHORS: Dylan Campbell, Lars Petersson, Laurent Kneip, Hongdong Li, Stephen Gould
HIGHLIGHT: Hence, we cast the problem as a 2D-3D mixture model alignment task and propose the first globally-optimal solution to this formulation under the robust L2 distance between mixture distributions.