

TITLE: 3D Point Capsule Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_3D_Point_Capsule_Networks_CVPR_2019_paper.html
AUTHORS: Yongheng Zhao, Tolga Birdal, Haowen Deng, Federico Tombari
HIGHLIGHT: In this paper, we propose 3D point-capsule networks, an auto-encoder designed to process sparse 3D point clouds while preserving spatial arrangements of the input data.

TITLE: GS3D: An Efficient 3D Object Detection Framework for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_GS3D_An_Efficient_3D_Object_Detection_Framework_for_Autonomous_Driving_CVPR_2019_paper.html
AUTHORS: Buyu Li, Wanli Ouyang, Lu Sheng, Xingyu Zeng, Xiaogang Wang
HIGHLIGHT: We present an efficient 3D object detection framework based on a single RGB image in the scenario of autonomous driving.

TITLE: Single-Image Piece-Wise Planar 3D Reconstruction via Associative Embedding
http://openaccess.thecvf.com/content_CVPR_2019/html/Yu_Single-Image_Piece-Wise_Planar_3D_Reconstruction_via_Associative_Embedding_CVPR_2019_paper.html
AUTHORS: Zehao Yu, Jia Zheng, Dongze Lian, Zihan Zhou, Shenghua Gao
HIGHLIGHT: To tackle this problem, we propose a novel two-stage method based on associative embedding, inspired by its recent success in instance segmentation.

TITLE: 3DN: 3D Deformation Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_3DN_3D_Deformation_Network_CVPR_2019_paper.html
AUTHORS: Weiyue Wang, Duygu Ceylan, Radomir Mech, Ulrich Neumann
HIGHLIGHT: Given such a source 3D model and a target which can be a 2D image, 3D model, or a point cloud acquired as a depth scan, we introduce 3DN, an end-to-end network that deforms the source model to resemble the target.

TITLE: HorizonNet: Learning Room Layout With 1D Representation and Pano Stretch Data Augmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_HorizonNet_Learning_Room_Layout_With_1D_Representation_and_Pano_Stretch_CVPR_2019_paper.html
AUTHORS: Cheng Sun, Chi-Wei Hsiao, Min Sun, Hwann-Tzong Chen
HIGHLIGHT: We present a new approach to the problem of estimating the 3D room layout from a single panoramic image.

TITLE: Deep Fitting Degree Scoring Network for Monocular 3D Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Deep_Fitting_Degree_Scoring_Network_for_Monocular_3D_Object_Detection_CVPR_2019_paper.html
AUTHORS: Lijie Liu, Jiwen Lu, Chunjing Xu, Qi Tian, Jie Zhou
HIGHLIGHT: In this paper, we propose to learn a deep fitting degree scoring network for monocular 3D object detection, which aims to score fitting degree between proposals and object conclusively.

TITLE: Dense Depth Posterior (DDP) From Single Image and Sparse Range
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Dense_Depth_Posterior_DDP_From_Single_Image_and_Sparse_Range_CVPR_2019_paper.html
AUTHORS: Yanchao Yang, Alex Wong, Stefano Soatto
HIGHLIGHT: We present a deep learning system to infer the posterior distribution of a dense depth map associated with an image, by exploiting sparse range measurements, for instance from a lidar.

TITLE: DuLa-Net: A Dual-Projection Network for Estimating Room Layouts From a Single RGB Panorama
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_DuLa-Net_A_Dual-Projection_Network_for_Estimating_Room_Layouts_From_a_CVPR_2019_paper.html
AUTHORS: Shang-Ta Yang, Fu-En Wang, Chi-Han Peng, Peter Wonka, Min Sun, Hung-Kuo Chu
HIGHLIGHT: We present a deep learning framework, called DuLa-Net, to predict Manhattan-world 3D room layouts from a single RGB panorama.

TITLE: Veritatem Dies Aperit - Temporally Consistent Depth Prediction Enabled by a Multi-Task Geometric and Semantic Scene Understanding Approach
http://openaccess.thecvf.com/content_CVPR_2019/html/Atapour-Abarghouei_Veritatem_Dies_Aperit_-_Temporally_Consistent_Depth_Prediction_Enabled_by_CVPR_2019_paper.html
AUTHORS: Amir Atapour-Abarghouei, Toby P. Breckon
HIGHLIGHT: In this paper, we propose a multi-task learning-based approach capable of jointly performing geometric and semantic scene understanding, namely depth prediction (monocular depth estimation and depth completion) and semantic scene segmentation.

TITLE: Segmentation-Driven 6D Object Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Hu_Segmentation-Driven_6D_Object_Pose_Estimation_CVPR_2019_paper.html
AUTHORS: Yinlin Hu, Joachim Hugonot, Pascal Fua, Mathieu Salzmann
HIGHLIGHT: In this paper, we introduce a segmentation-driven 6D pose estimation framework where each visible part of the objects contributes a local pose prediction in the form of 2D keypoint locations.

TITLE: Exploiting Temporal Context for 3D Human Pose Estimation in the Wild
http://openaccess.thecvf.com/content_CVPR_2019/html/Arnab_Exploiting_Temporal_Context_for_3D_Human_Pose_Estimation_in_the_CVPR_2019_paper.html
AUTHORS: Anurag Arnab, Carl Doersch, Andrew Zisserman
HIGHLIGHT: We present a bundle-adjustment-based algorithm for recovering accurate 3D human pose and meshes from monocular videos.

TITLE: What Do Single-View 3D Reconstruction Networks Learn?
http://openaccess.thecvf.com/content_CVPR_2019/html/Tatarchenko_What_Do_Single-View_3D_Reconstruction_Networks_Learn_CVPR_2019_paper.html
AUTHORS: Maxim Tatarchenko, Stephan R. Richter, Rene Ranftl, Zhuwen Li, Vladlen Koltun, Thomas Brox
HIGHLIGHT: Convolutional networks for single-view object reconstruction have shown impressive performance and have become a popular subject of research.

TITLE: 3D-SIS: 3D Semantic Instance Segmentation of RGB-D Scans
http://openaccess.thecvf.com/content_CVPR_2019/html/Hou_3D-SIS_3D_Semantic_Instance_Segmentation_of_RGB-D_Scans_CVPR_2019_paper.html
AUTHORS: Ji Hou, Angela Dai, Matthias Niessner
HIGHLIGHT: We introduce 3D-SIS, a novel neural network architecture for 3D semantic instance segmentation in commodity RGB-D scans.

TITLE: Causes and Corrections for Bimodal Multi-Path Scanning With Structured Light
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Causes_and_Corrections_for_Bimodal_Multi-Path_Scanning_With_Structured_Light_CVPR_2019_paper.html
AUTHORS: Yu Zhang, Daniel L. Lau, Ying Yu
HIGHLIGHT: In this paper, we present a general mathematical model to address the bimodal multi-path issue in a phase-measuring-profilometry scanner to measure the constructive and destructive interference between the two light paths, and by taking advantage of this interesting cue, separate the paths and make two decoupled phase measurements.

TITLE: TextureNet: Consistent Local Parametrizations for Learning From High-Resolution Signals on Meshes
http://openaccess.thecvf.com/content_CVPR_2019/html/Huang_TextureNet_Consistent_Local_Parametrizations_for_Learning_From_High-Resolution_Signals_on_CVPR_2019_paper.html
AUTHORS: Jingwei Huang, Haotian Zhang, Li Yi, Thomas Funkhouser, Matthias Niessner, Leonidas J. Guibas
HIGHLIGHT: We introduce, TextureNet, a neural network architecture designed to extract features from high-resolution signals associated with 3D surface meshes (e.g., color texture maps).

TITLE: PlaneRCNN: 3D Plane Detection and Reconstruction From a Single Image
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_PlaneRCNN_3D_Plane_Detection_and_Reconstruction_From_a_Single_Image_CVPR_2019_paper.html
AUTHORS: Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, Jan Kautz
HIGHLIGHT: This paper proposes a deep neural architecture, PlaneRCNN, that detects and reconstructs piecewise planar regions from a single RGB image.

TITLE: Occupancy Networks: Learning 3D Reconstruction in Function Space
http://openaccess.thecvf.com/content_CVPR_2019/html/Mescheder_Occupancy_Networks_Learning_3D_Reconstruction_in_Function_Space_CVPR_2019_paper.html
AUTHORS: Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, Andreas Geiger
HIGHLIGHT: In this paper, we propose Occupancy Networks, a new representation for learning-based 3D reconstruction methods.

TITLE: 3D Shape Reconstruction From Images in the Frequency Domain
http://openaccess.thecvf.com/content_CVPR_2019/html/Shen_3D_Shape_Reconstruction_From_Images_in_the_Frequency_Domain_CVPR_2019_paper.html
AUTHORS: Weichao Shen, Yunde Jia, Yuwei Wu

HIGHLIGHT: In this paper, we propose a Fourier-based method that reconstructs a 3D shape from images in a 2D space by predicting slices in the frequency domain.

TITLE: SiCloPe: Silhouette-Based Clothed People

http://openaccess.thecvf.com/content_CVPR_2019/html/Natsume_SiCloPe_Silhouette-Based_Clothed_People_CVPR_2019_paper.html

AUTHORS: Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, Shigeo Morishima

HIGHLIGHT: We introduce a new silhouette-based representation for modeling clothed human bodies using deep generative models.

TITLE: Detailed Human Shape Estimation From a Single Image by Hierarchical Mesh Deformation

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhu_Detailed_Human_Shape_Estimation_From_a_Single_Image_by_Hierarchical_CVPR_2019_paper.html

AUTHORS: Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, Ruigang Yang

HIGHLIGHT: This paper presents a novel framework to recover detailed human body shapes from a single image.

TITLE: Convolutional Mesh Regression for Single-Image Human Shape Reconstruction

http://openaccess.thecvf.com/content_CVPR_2019/html/Kolotouros_Convolutional_Mesh_Regression_for_Single-Image_Human_Shape_Reconstruction_CVPR_2019_paper.html

AUTHORS: Nikos Kolotouros, Georgios Pavlakos, Kostas Daniilidis

HIGHLIGHT: In our work, we propose to relax this heavy reliance on the model's parameter space.

TITLE: H+O: Unified Egocentric Recognition of 3D Hand-Object Poses and Interactions

http://openaccess.thecvf.com/content_CVPR_2019/html/Tekin_HO_Unified_Egocentric_Recognition_of_3D_Hand-Object_Poses_and_Interactions_CVPR_2019_paper.html

AUTHORS: Bugra Tekin, Federica Bogo, Marc Pollefeys

HIGHLIGHT: We present a unified framework for understanding 3D hand and object interactions in raw image sequences from egocentric RGB cameras.

TITLE: Learning the Depths of Moving People by Watching Frozen People

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Learning_the_Depths_of_Moving_People_by_Watching_Frozen_People_CVPR_2019_paper.html

AUTHORS: Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, William T. Freeman

HIGHLIGHT: We present a method for predicting dense depth in scenarios where both a monocular camera and people in the scene are freely moving.

TITLE: Extreme Relative Pose Estimation for RGB-D Scans via Scene Completion

http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Extreme_Relative_Pose_Estimation_for_RGB-D_Scans_via_Scene_Completion_CVPR_2019_paper.html

AUTHORS: Zhenpei Yang, Jeffrey Z. Pan, Linjie Luo, Xiaowei Zhou, Kristen Grauman, Qixing Huang

HIGHLIGHT: We introduce a novel approach that extends the scope to extreme relative poses, with little or even no overlap between the input scans.

TITLE: A Skeleton-Bridged Deep Learning Approach for Generating Meshes of Complex Topologies From Single RGB Images

http://openaccess.thecvf.com/content_CVPR_2019/html/Tang_A_Skeleton-Bridged_Deep_Learning_Approach_for_Generating_Meshes_of_Complex_Topologies_From_Single_RGB_Images_CVPR_2019_paper.html

AUTHORS: Jiapeng Tang, Xiaoguang Han, Junyi Pan, Kui Jia, Xin Tong

HIGHLIGHT: To this end, we propose in this paper a skeleton-bridged, stage-wise learning approach to address the challenge.

TITLE: Learning Structure-And-Motion-Aware Rolling Shutter Correction

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhuang_Learning_Structure-And-Motion-Aware_Rolling_Shutter_Correction_CVPR_2019_paper.html

AUTHORS: Bingbing Zhuang, Quoc-Huy Tran, Pan Ji, Loong-Fah Cheong, Manmohan Chandraker

HIGHLIGHT: Our method learns from a large-scale dataset synthesized in a geometrically meaningful way where the RS effect is generated in a manner consistent with the camera motion and scene structure.

TITLE: PVNet: Pixel-Wise Voting Network for 6DoF Pose Estimation

http://openaccess.thecvf.com/content_CVPR_2019/html/Peng_PVNet_Pixel-Wise_Voting_Network_for_6DoF_Pose_Estimation_CVPR_2019_paper.html

AUTHORS: Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, Hujun Bao

HIGHLIGHT: Instead, we introduce a Pixel-wise Voting Network (PVNet) to regress pixel-wise vectors pointing to the keypoints and use these vectors to vote for keypoint locations.

TITLE: Learning Single-Image Depth From Videos Using Quality Assessment Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Learning_Single-Image_Depth_From_Videos_Using_Quality_Assessment_Networks_CVPR_2019_paper.html

AUTHORS: Weifeng Chen, Shengyi Qian, Jia Deng

HIGHLIGHT: In this paper we propose a method to automatically generate such data through Structure-from-Motion (SfM) on Internet videos.

TITLE: Learning 3D Human Dynamics From Video
http://openaccess.thecvf.com/content_CVPR_2019/html/Kanazawa_Learning_3D_Human_Dynamics_From_Video_CVPR_2019_paper.html

AUTHORS: Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, Jitendra Malik

HIGHLIGHT: We present a framework that can similarly learn a representation of 3D dynamics of humans from video via a simple but effective temporal encoding of image features.

TITLE: Lending Orientation to Neural Networks for Cross-View Geo-Localization
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Lending_Orientation_to_Neural_Networks_for_Cross-View_Geo-Localization_CVPR_2019_paper.html

AUTHORS: Liu Liu, Hongdong Li

HIGHLIGHT: Inspired by this insight, this paper proposes a novel method which endows deep neural networks with the 'commonsense' of orientation.

TITLE: Visual Localization by Learning Objects-Of-Interest Dense Match Regression
http://openaccess.thecvf.com/content_CVPR_2019/html/Weinzaepfel_Visual_Localization_by_Learning_Objects-Of-Interest_Dense_Match_Regression_CVPR_2019_paper.html

AUTHORS: Philippe Weinzaepfel, Gabriela Csurka, Yohann Cabon, Martin Humenberger

HIGHLIGHT: We introduce a novel CNN-based approach for visual localization from a single RGB image that relies on densely matching a set of Objects-of-Interest (OOIs).

TITLE: Bilateral Cyclic Constraint and Adaptive Regularization for Unsupervised Monocular Depth Prediction
http://openaccess.thecvf.com/content_CVPR_2019/html/Wong_Bilateral_Cyclic_Constraint_and_Adaptive_Regularization_for_Unsupervised_Monocular_Depth_CVPR_2019_paper.html

AUTHORS: Alex Wong, Stefano Soatto

HIGHLIGHT: We propose a novel objective function that exploits the bilateral cyclic relationship between the left and right disparities and we introduce an adaptive regularization scheme that allows the network to handle both the co-visible and occluded regions in a stereo pair.

TITLE: MMFace: A Multi-Metric Regression Network for Unconstrained Face Reconstruction
http://openaccess.thecvf.com/content_CVPR_2019/html/Yi_MMFace_A_Multi-Metric_Regression_Network_for_Unconstrained_Face_Reconstruction_CVPR_2019_paper.html

AUTHORS: Hongwei Yi, Chen Li, Qiong Cao, Xiaoyong Shen, Sheng Li, Guoping Wang, Yu-Wing Tai

HIGHLIGHT: We propose to address the face reconstruction in the wild by using a multi-metric regression network, MMFace, to align a 3D face morphable model (3DMM) to an input image.

TITLE: 3D Motion Decomposition for RGBD Future Dynamic Scene Synthesis
http://openaccess.thecvf.com/content_CVPR_2019/html/Qi_3D_Motion_Decomposition_for_RGBD_Future_Dynamic_Scene_Synthesis_CVPR_2019_paper.html

AUTHORS: Xiaojuan Qi, Zhengzhe Liu, Qifeng Chen, Jiaya Jia

HIGHLIGHT: In this paper, we propose a RGBD scene forecasting model with 3D motion decomposition.

TITLE: Single Image Depth Estimation Trained via Depth From Defocus Cues
http://openaccess.thecvf.com/content_CVPR_2019/html/Gur_Single_Image_Depth_Estimation_Trained_via_Depth_From_Defocus_Cues_CVPR_2019_paper.html

AUTHORS: Shir Gur, Lior Wolf

HIGHLIGHT: In this work, we rely, instead of different views, on depth from focus cues.

TITLE: RGBD Based Dimensional Decomposition Residual Network for 3D Semantic Scene Completion
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_RGBD_Based_Dimensional_Decomposition_Residual_Network_for_3D_Semantic_Scene_CVPR_2019_paper.html

AUTHORS: Jie Li, Yu Liu, Dong Gong, Qinfeng Shi, Xia Yuan, Chunxia Zhao, Ian Reid
HIGHLIGHT: We introduce a light-weight Dimensional Decomposition Residual network (DDR) for 3D dense prediction tasks.

TITLE: Neural Scene Decomposition for Multi-Person Motion Capture
http://openaccess.thecvf.com/content_CVPR_2019/html/Rhodin_Neural_Scene_Decomposition_for_Multi-Person_Motion_Capture_CVPR_2019_paper.html

AUTHORS: Helge Rhodin, Victor Constantin, Isinsu Katircioglu, Mathieu Salzmann, Pascal Fua
HIGHLIGHT: In this paper, we therefore propose an approach to learning features that are useful for this purpose.

TITLE: Minimal Solvers for Mini-Loop Closures in 3D Multi-Scan Alignment
http://openaccess.thecvf.com/content_CVPR_2019/html/Miraldo_Minimal_Solvers_for_Mini-Loop_Closures_in_3D_Multi-Scan_Alignment_CVPR_2019_paper.html

AUTHORS: Pedro Miraldo, Surojit Saha, Srikumar Ramalingam
HIGHLIGHT: In this paper, we take a different approach and develop minimal solvers for jointly computing the initial poses of cameras in small loops such as 3-, 4-, and 5-cycles.

TITLE: Volumetric Capture of Humans With a Single RGBD Camera via Semi-Parametric Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Pandey_Volumetric_Capture_of_Humans_With_a_Single_RGBD_Camera_via_CVPR_2019_paper.html

AUTHORS: Rohit Pandey, Anastasia Tkach, Shuoran Yang, Pavel Pidlypenskyi, Jonathan Taylor, Ricardo Martin-Brualla, Andrea Tagliasacchi, George Papandreou, Philip Davidson, Cem Keskin, Shahrnam Izadi, Sean Fanello
HIGHLIGHT: Thus, in this work, we propose a method to synthesize free viewpoint renderings using a single RGBD camera.

TITLE: Joint Face Detection and Facial Motion Retargeting for Multiple Faces
http://openaccess.thecvf.com/content_CVPR_2019/html/Chaudhuri_Joint_Face_Detection_and_Facial_Motion_Retargeting_for_Multiple_Faces_CVPR_2019_paper.html

AUTHORS: Bindita Chaudhuri, Noranart Vesdapunt, Baoyuan Wang
HIGHLIGHT: In this paper, we present a single end-to-end network to jointly predict the bounding box locations and 3DMM parameters for multiple faces.

TITLE: Monocular Depth Estimation Using Relative Depth Maps
http://openaccess.thecvf.com/content_CVPR_2019/html/Lee_Monocular_Depth_Estimation_Using_Relative_Depth_Maps_CVPR_2019_paper.html

AUTHORS: Jae-Han Lee, Chang-Su Kim
HIGHLIGHT: We propose a novel algorithm for monocular depth estimation using relative depth maps.

TITLE: Unsupervised Primitive Discovery for Improved 3D Generative Modeling
http://openaccess.thecvf.com/content_CVPR_2019/html/Khan_Unsupervised_Primitive_Discovery_for_Improved_3D_Generative_Modeling_CVPR_2019_paper.html

AUTHORS: Salman H. Khan, Yulan Guo, Munawar Hayat, Nick Barnes
HIGHLIGHT: Here, we propose a novel factorized generative model for 3D shape generation that sequentially transitions from coarse to fine scale shape generation.

TITLE: Learning to Explore Intrinsic Saliency for Stereoscopic Video
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Learning_to_Explore_Intrinsic_Saliency_for_Stereoscopic_Video_CVPR_2019_paper.html

AUTHORS: Qiudan Zhang, Xu Wang, Shiqi Wang, Shikai Li, Sam Kwong, Jianmin Jiang
HIGHLIGHT: In this paper, we argue that the high-level features are crucial and resort to the deep learning framework to learn the saliency map of stereoscopic videos.

TITLE: Spherical Regression: Learning Viewpoints, Surface Normals and 3D Rotations on N-Spheres
http://openaccess.thecvf.com/content_CVPR_2019/html/Liao_Spherical_Regression_Learning_Viewpoints_Surface_Normals_and_3D_Rotations_on_CVPR_2019_paper.html

AUTHORS: Shuai Liao, Efstratios Gavves, Cees G. M. Snoek
HIGHLIGHT: We observe many continuous output problems in computer vision are naturally contained in closed geometrical manifolds, like the Euler angles in viewpoint estimation or the normals in surface normal estimation.

TITLE: Refine and Distill: Exploiting Cycle-Inconsistency and Knowledge Distillation for Unsupervised Monocular Depth Estimation

http://openaccess.thecvf.com/content_CVPR_2019/html/Pilzer_Refine_and_Distill_Exploiting_Cycle-Inconsistency_and_Knowledge_Distillation_for_Unsupervised_CVPR_2019_paper.html

AUTHORS: Andrea Pilzer, Stephane Lathuiliere, Nicu Sebe, Elisa Ricci

HIGHLIGHT: Following these works, we propose a novel self-supervised deep model for estimating depth maps.

TITLE: Learning View Priors for Single-View 3D Reconstruction

http://openaccess.thecvf.com/content_CVPR_2019/html/Kato_Learning_View_Priors_for_Single-View_3D_Reconstruction_CVPR_2019_paper.html

AUTHORS: Hiroharu Kato, Tatsuya Harada

HIGHLIGHT: To reconstruct shapes that look reasonable from any viewpoint, we propose to train a discriminator that learns prior knowledge regarding possible views.

TITLE: Geometry-Aware Symmetric Domain Adaptation for Monocular Depth Estimation

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Geometry-Aware_Symmetric_Domain_Adaptation_for_Monocular_Depth_Estimation_CVPR_2019_paper.html

AUTHORS: Shanshan Zhao, Huan Fu, Mingming Gong, Dacheng Tao

HIGHLIGHT: Motivated by the observation, we propose a geometry-aware symmetric domain adaptation framework (GASDA) to explore the labels in the synthetic data and epipolar geometry in the real data jointly.

TITLE: Learning Monocular Depth Estimation Infusing Traditional Stereo Knowledge

http://openaccess.thecvf.com/content_CVPR_2019/html/Tosi_Learning_Monocular_Depth_Estimation_Infusing_Traditional_Stereo_Knowledge_CVPR_2019_paper.html

AUTHORS: Fabio Tosi, Filippo Aleotti, Matteo Poggi, Stefano Mattocchia

HIGHLIGHT: To this aim we propose monoResMatch, a novel deep architecture designed to infer depth from a single input image by synthesizing features from a different point of view, horizontally aligned with the input image, performing stereo matching between the two cues.

TITLE: SIGNet: Semantic Instance Aided Unsupervised 3D Geometry Perception

http://openaccess.thecvf.com/content_CVPR_2019/html/Meng_SIGNet_Semantic_Instance_Aided_Unsupervised_3D_Geometry_Perception_CVPR_2019_paper.html

AUTHORS: Yue Meng, Yongxi Lu, Aman Raj, Samuel Sunarjo, Rui Guo, Tara Javidi, Gaurav Bansal, Dinesh Bharadia

HIGHLIGHT: This paper introduces SIGNet, a novel framework that provides robust geometry perception without requiring geometrically informative labels.

TITLE: Learning Joint Reconstruction of Hands and Manipulated Objects

http://openaccess.thecvf.com/content_CVPR_2019/html/Hasson_Learning_Joint_Reconstruction_of_Hands_and_Manipulated_Objects_CVPR_2019_paper.html

AUTHORS: Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, Cordelia Schmid

HIGHLIGHT: In this work, we regularize the joint reconstruction of hands and objects with manipulation constraints.

TITLE: Deep Single Image Camera Calibration With Radial Distortion

http://openaccess.thecvf.com/content_CVPR_2019/html/Lopez_Deep_Single_Image_Camera_Calibration_With_Radial_Distortion_CVPR_2019_paper.html

AUTHORS: Manuel Lopez, Roger Mari, Pau Gargallo, Yubin Kuang, Javier Gonzalez-Jimenez, Gloria Haro

HIGHLIGHT: In this work we propose a method to predict extrinsic (tilt and roll) and intrinsic (focal length and radial distortion) parameters from a single image.

TITLE: CAM-ConvS: Camera-Aware Multi-Scale Convolutions for Single-View Depth

http://openaccess.thecvf.com/content_CVPR_2019/html/Facil_CAM-ConvS_Camera-Aware_Multi-Scale_Convolutions_for_Single-View_Depth_CVPR_2019_paper.html

AUTHORS: Jose M. Facil, Benjamin Ummenhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, Javier Civera

HIGHLIGHT: In this work, we propose a new type of convolution that can take the camera parameters into account, thus allowing neural networks to learn calibration-aware patterns.

TITLE: Translate-to-Recognize Networks for RGB-D Scene Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Du_Translate-to-Recognize_Networks_for_RGB-D_Scene_Recognition_CVPR_2019_paper.html

AUTHORS: Dapeng Du, Limin Wang, Huiling Wang, Kai Zhao, Gangshan Wu

HIGHLIGHT: To this end, this paper presents a unified framework to integrate the tasks of cross-modal translation and modality-specific recognition, termed as Translate-to-Recognize Network TRecgNet.

TITLE: Re-Identification Supervised Texture Generation
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Re-Identification_Supervised_Texture_Generation_CVPR_2019_paper.html
AUTHORS: Jian Wang, Yunshan Zhong, Yachun Li, Chi Zhang, Yichen Wei
HIGHLIGHT: In this paper, we propose an end-to-end learning strategy to generate textures of human bodies under the supervision of person re-identification.

TITLE: Action4D: Online Action Recognition in the Crowd and Clutter
http://openaccess.thecvf.com/content_CVPR_2019/html/You_Action4D_Online_Action_Recognition_in_the_Crowd_and_Clutter_CVPR_2019_paper.html
AUTHORS: Quanzeng You, Hao Jiang
HIGHLIGHT: At the first step, we propose a new method to track people in 4D, which can reliably detect and follow each person in real time.

TITLE: Monocular 3D Object Detection Leveraging Accurate Proposals and Shape Reconstruction
http://openaccess.thecvf.com/content_CVPR_2019/html/Ku_Monocular_3D_Object_Detection_Leveraging_Accurate_Proposals_and_Shape_Reconstruction_CVPR_2019_paper.html
AUTHORS: Jason Ku, Alex D. Pon, Steven L. Waslander
HIGHLIGHT: We present MonoPSR, a monocular 3D object detection method that leverages proposals and shape reconstruction.