

TITLE: Reducing Uncertainty in Undersampled MRI Reconstruction With Active Acquisition
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Reducing_Uncertainty_in_Undersampled_MRI_Reconstruction_With_Active_Acquisition_CVPR_2019_paper.html
AUTHORS: Zizhao Zhang, Adriana Romero, Matthew J. Muckley, Pascal Vincent, Lin Yang, Michal Drozdal
HIGHLIGHT: In this paper, we present a novel method for MRI reconstruction that, at inference time, dynamically selects the measurements to take and iteratively refines the prediction in order to best reduce the reconstruction error and, thus, its uncertainty.

TITLE: ESIR: End-To-End Scene Text Recognition via Iterative Image Rectification
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhan_ESIR_End-To-End_Scene_Text_Recognition_via_Iterative_Image_Rectification_CVPR_2019_paper.html
AUTHORS: Fangneng Zhan, Shijian Lu
HIGHLIGHT: This paper presents an end-to-end trainable scene text recognition system (ESIR) that iteratively removes perspective distortion and text line curvature as driven by better scene text recognition performance.

TITLE: ROI-10D: Monocular Lifting of 2D Detection to 6D Pose and Metric Shape
http://openaccess.thecvf.com/content_CVPR_2019/html/Manhardt_ROI-10D_Monocular_Lifting_of_2D_Detection_to_6D_Pose_and_CVPR_2019_paper.html
AUTHORS: Fabian Manhardt, Wadim Kehl, Adrien Gaidon
HIGHLIGHT: We present a deep learning method for end-to-end monocular 3D object detection and metric shape retrieval.

TITLE: Collaborative Learning of Semi-Supervised Segmentation and Classification for Medical Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhou_Collaborative_Learning_of_Semi-Supervised_Segmentation_and_Classification_for_Medical_Images_CVPR_2019_paper.html
AUTHORS: Yi Zhou, Xiaodong He, Lei Huang, Li Liu, Fan Zhu, Shanshan Cui, Ling Shao
HIGHLIGHT: In this paper, we propose a collaborative learning method to jointly improve the performance of disease grading and lesion segmentation by semi-supervised learning with an attention mechanism.

TITLE: Biologically-Constrained Graphs for Global Connectomics Reconstruction
http://openaccess.thecvf.com/content_CVPR_2019/html/Matejek_Biologically-Constrained_Graphs_for_Global_Connectomics_Reconstruction_CVPR_2019_paper.html
AUTHORS: Brian Matejek, Daniel Haehn, Haidong Zhu, Donglai Wei, Toufiq Parag, Hanspeter Pfister
HIGHLIGHT: We propose a third step for connectomics reconstruction pipelines to refine an over-segmentation using both local and global context with an emphasis on adhering to the underlying biology.

TITLE: P3SGD: Patient Privacy Preserving SGD for Regularizing Deep CNNs in Pathological Image Classification
http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_P3SGD_Patient_Privacy_Preserving_SGD_for_Regularizing_Deep_CN_Ns_in_CVPR_2019_paper.html
AUTHORS: Bingzhe Wu, Shiwan Zhao, Guangyu Sun, Xiaolu Zhang, Zhong Su, Caihong Zeng, Zhihong Liu
HIGHLIGHT: To tackle the above two challenges, we introduce a novel stochastic gradient descent (SGD) scheme, named patient privacy preserving SGD (P3SGD), which performs the model update of the SGD in the patient level via a large-step update built upon each patient's data.

TITLE: Elastic Boundary Projection for 3D Medical Image Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Ni_Elastic_Boundary_Projection_for_3D_Medical_Image_Segmentation_CVPR_2019_paper.html
AUTHORS: Tianwei Ni, Lingxi Xie, Huangjie Zheng, Elliot K. Fishman, Alan L. Yuille
HIGHLIGHT: In this paper, we bridge the gap between 2D and 3D using a novel approach named Elastic Boundary Projection (EBP).

TITLE: SIXray: A Large-Scale Security Inspection X-Ray Benchmark for Prohibited Item Discovery in Overlapping Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Miao_SIXray_A_Large-Scale_Security_Inspection_X-Ray_Benchmark_for_Prohibited_Item_CVPR_2019_paper.html
AUTHORS: Caijing Miao, Lingxi Xie, Fang Wan, Chi Su, Hongye Liu, Jianbin Jiao, Qixiang Ye
HIGHLIGHT: In this paper, we present a large-scale dataset and establish a baseline for prohibited item discovery in Security Inspection X-ray images.

TITLE: Noise2Void - Learning Denoising From Single Noisy Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Krull_Noise2Void_-_Learning_Denoising_From_Single_Noisy_Images_CVPR_2019_paper.html
AUTHORS: Alexander Krull, Tim-Oliver Buchholz, Florian Jug
HIGHLIGHT: Here, we introduce Noise2Void (N2V), a training scheme that takes this idea one step further.

TITLE: CRAVES: Controlling Robotic Arm With a Vision-Based Economic System
http://openaccess.thecvf.com/content_CVPR_2019/html/Zuo_CRAVES_Controlling_Robotic_Arm_With_a_Vision-Based_Economic_System_CVPR_2019_paper.html
AUTHORS: Yiming Zuo, Weichao Qiu, Lingxi Xie, Fangwei Zhong, Yizhou Wang, Alan L. Yuille
HIGHLIGHT: In this paper, we present an alternative solution, which uses a 3D model to create a large number of synthetic data, trains a vision model in this virtual domain, and applies it to real-world images after domain adaptation.

TITLE: Networks for Joint Affine and Non-Parametric Image Registration
http://openaccess.thecvf.com/content_CVPR_2019/html/Shen_Networks_for_Joint_Affine_and_Non-Parametric_Image_Registration_CVPR_2019_paper.html
AUTHORS: Zhengyang Shen, Xu Han, Zhenlin Xu, Marc Niethammer
HIGHLIGHT: We introduce an end-to-end deep-learning framework for 3D medical image registration.

TITLE: Learning Shape-Aware Embedding for Scene Text Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Tian_Learning_Shape-Aware_Embedding_for_Scene_Text_Detection_CVPR_2019_paper.html
AUTHORS: Zhuotao Tian, Michelle Shu, Pengyuan Lyu, Ruiyu Li, Chao Zhou, Xiaoyong Shen, Jiaya Jia
HIGHLIGHT: Specifically, we treat text detection as instance segmentation and propose a segmentation-based framework, which extracts each text instance as an independent connected component.

TITLE: Learning to Film From Professional Human Motion Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Huang_Learning_to_Film_From_Professional_Human_Motion_Videos_CVPR_2019_paper.html
AUTHORS: Chong Huang, Chuan-En Lin, Zhenyu Yang, Yan Kong, Peng Chen, Xin Yang, Kwang-Ting Cheng
HIGHLIGHT: In this study, we propose a learning-based framework which incorporates the video contents and previous camera motions to predict the future camera motions that enable the capture of professional videos.

TITLE: Pay Attention! - Robustifying a Deep Visuomotor Policy Through Task-Focused Visual Attention
http://openaccess.thecvf.com/content_CVPR_2019/html/Abolghasemi_Pay_Attention_-_Robustifying_a_Deep_Visuomotor_Policy_Through_Task-Focused_CVPR_2019_paper.html
AUTHORS: Pooya Abolghasemi, Amir Mazaheri, Mubarak Shah, Ladislau Boloni
HIGHLIGHT: In this paper, we propose an approach for augmenting a deep visuomotor policy trained through demonstrations with Task Focused visual Attention (TFA).

TITLE: Deep Blind Video Decaptioning by Temporal Aggregation and Recurrence
http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Deep_Blind_Video_Decaptioning_by_Temporal_Aggregation_and_Recurrence_CVPR_2019_paper.html
AUTHORS: Dahun Kim, Sanghyun Woo, Joon-Young Lee, In So Kweon
HIGHLIGHT: In this paper, we propose a simple yet effective framework for fast blind video decaptioning.

TITLE: Noise-Aware Unsupervised Deep Lidar-Stereo Fusion
http://openaccess.thecvf.com/content_CVPR_2019/html/Cheng_Noise-Aware_Unsupervised_Deep_Lidar-Stereo_Fusion_CVPR_2019_paper.html
AUTHORS: Xuelian Cheng, Yiran Zhong, Yuchao Dai, Pan Ji, Hongdong Li
HIGHLIGHT: In this paper, we present LidarStereoNet, the first unsupervised Lidar-stereo fusion network, which can be trained in an end-to-end manner without the need of ground truth depth maps.

TITLE: Versatile Multiple Choice Learning and Its Application to Vision Computing
http://openaccess.thecvf.com/content_CVPR_2019/html/Tian_Versatile_Multiple_Choice_Learning_and_Its_Application_to_Vision_Computing_CVPR_2019_paper.html
AUTHORS: Kai Tian, Yi Xu, Shuigeng Zhou, Jihong Guan
HIGHLIGHT: In this paper, a new MCL method, called vMCL (the abbreviation of versatile Multiple Choice Learning), is developed to extend the application scenarios of MCL methods by ensembling deep neural networks.

TITLE: EV-Gait: Event-Based Robust Gait Recognition Using Dynamic Vision Sensors
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_EV-Gait_Event-Based_Robust_Gait_Recognition_Using_Dynamic_Vision_Sensors_CVPR_2019_paper.html
AUTHORS: Yanxiang Wang, Bowen Du, Yiran Shen, Kai Wu, Guangrong Zhao, Jianguo Sun, Hongkai Wen
HIGHLIGHT: In this paper, we introduce a new type of sensing modality, the Dynamic Vision Sensors (Event Cameras), for the task of gait recognition.

TITLE: ToothNet: Automatic Tooth Instance Segmentation and Identification From Cone Beam CT Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Cui_ToothNet_Automatic_Tooth_Instance_Segmentation_and_Identification_From_Cone_Beam_CVPR_2019_paper.html
AUTHORS: Zhiming Cui, Changjian Li, Wenping Wang
HIGHLIGHT: This paper proposes a method that uses deep convolutional neural networks to achieve automatic and accurate tooth instance segmentation and identification from CBCT (cone beam CT) images for digital dentistry.

TITLE: Modularized Textual Grounding for Counterfactual Resilience
http://openaccess.thecvf.com/content_CVPR_2019/html/Fang_Modularized_Textual_Grounding_for_Counterfactual_Resilience_CVPR_2019_paper.html
AUTHORS: Zhiyuan Fang, Shu Kong, Charless Fowlkes, Yezhou Yang
HIGHLIGHT: To address these issues, we propose a visual grounding system which is 1) end-to-end trainable in a weakly supervised fashion with only image-level annotations, and 2) counterfactually resilient owing to the modular design.

TITLE: L3-Net: Towards Learning Based LiDAR Localization for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2019/html/Lu_L3-Net_Towards_Learning_Based_LiDAR_Localization_for_Autonomous_Driving_CVPR_2019_paper.html
AUTHORS: Weixin Lu, Yao Zhou, Guowei Wan, Shenhua Hou, Shiyu Song
HIGHLIGHT: We present L3-Net - a novel learning-based LiDAR localization system that achieves centimeter-level localization accuracy, comparable to prior state-of-the-art systems with hand-crafted pipelines.

TITLE: Uncertainty Guided Multi-Scale Residual Learning-Using a Cycle Spinning CNN for Single Image De-Raining
http://openaccess.thecvf.com/content_CVPR_2019/html/Yasarla_Uncertainty_Guided_Multi-Scale_Residual_Learning-Using_a_Cycle_Spinning_CNN_for_CVPR_2019_paper.html
AUTHORS: Rajeev Yasarla, Vishal M. Patel
HIGHLIGHT: Previous approaches have attempted to address this problem by leveraging some prior information to remove rain streaks from a single image.

TITLE: Toward Realistic Image Compositing With Adversarial Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Toward_Realistic_Image_Compositing_With_Adversarial_Learning_CVPR_2019_paper.html
AUTHORS: Bor-Chun Chen, Andrew Kae
HIGHLIGHT: In this work we propose a generative adversarial network (GAN) architecture for automatic image compositing.

TITLE: Cross-Classification Clustering: An Efficient Multi-Object Tracking Technique for 3-D Instance Segmentation in Connectomics
http://openaccess.thecvf.com/content_CVPR_2019/html/Meirovitch_Cross-Classification_Clustering_An_Efficient_Multi-Object_Tracking_Technique_for_3-D_Instance_CVPR_2019_paper.html
AUTHORS: Yaron Meirovitch, Lu Mi, Hayk Saribekyan, Alexander Matveev, David Rolnick, Nir Shavit
HIGHLIGHT: Here we introduce cross-classification clustering (3C), a technique that simultaneously tracks complex, interrelated objects in an image stack.

TITLE: Deep ChArUco: Dark ChArUco Marker Pose Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Hu_Deep_ChArUco_Dark_ChArUco_Marker_Pose_Estimation_CVPR_2019_paper.html
AUTHORS: Danying Hu, Daniel DeTone, Tomasz Malisiewicz
HIGHLIGHT: We present Deep ChArUco, a real-time pose estimation system which combines two custom deep networks, ChArUcoNet and RefineNet, with the Perspective-n-Point (PnP) algorithm to estimate the marker's 6DoF pose.

TITLE: Pseudo-LiDAR From Visual Depth Estimation: Bridging the Gap in 3D Object Detection for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Pseudo-LiDAR_From_Visual_Depth_Estimation_Bridging_the_Gap_in_3D_CVPR_2019_paper.html
AUTHORS: Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger
HIGHLIGHT: Taking the inner workings of convolutional neural networks into consideration, we propose to convert image-based depth maps to pseudo-LiDAR representations --- essentially mimicking the LiDAR signal.

TITLE: Rules of the Road: Predicting Driving Behavior With a Convolutional Model of Semantic Interactions
http://openaccess.thecvf.com/content_CVPR_2019/html/Hong_Rules_of_the_Road_Predicting_Driving_Behavior_With_a_Convolutional_CVPR_2019_paper.html
AUTHORS: Joey Hong, Benjamin Sapp, James Philbin

HIGHLIGHT: We present a unified representation which encodes such high-level semantic information in a spatial grid, allowing the use of deep convolutional models to fuse complex scene context.

TITLE: Metric Learning for Image Registration

http://openaccess.thecvf.com/content_CVPR_2019/html/Niethammer_Metric_Learning_for_Image_Registration_CVPR_2019_paper.html

AUTHORS: Marc Niethammer, Roland Kwitt, Francois-Xavier Vialard

HIGHLIGHT: Instead of learning the entire registration approach, we learn a spatially-adaptive regularizer within a registration model.

TITLE: LO-Net: Deep Real-Time Lidar Odometry

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_LO-Net_Deep_Real-Time_Lidar_Odometry_CVPR_2019_paper.html

AUTHORS: Qing Li, Shaoyang Chen, Cheng Wang, Xin Li, Chenglu Wen, Ming Cheng, Jonathan Li

HIGHLIGHT: We present a novel deep convolutional network pipeline, LO-Net, for real-time lidar odometry estimation.

TITLE: TraPHic: Trajectory Prediction in Dense and Heterogeneous Traffic Using Weighted Interactions

http://openaccess.thecvf.com/content_CVPR_2019/html/Chandra_TraPHic_Trajectory_Prediction_in_Dense_and_Heterogeneous_Traffic_Using_Weighted_CVPR_2019_paper.html

AUTHORS: Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, Dinesh Manocha

HIGHLIGHT: We present a new algorithm for predicting the near-term trajectories of road agents in dense traffic videos.

TITLE: World From Blur

http://openaccess.thecvf.com/content_CVPR_2019/html/Qiu_World_From_Blur_CVPR_2019_paper.html

AUTHORS: Jiayan Qiu, Xinchao Wang, Stephen J. Maybank, Dacheng Tao

HIGHLIGHT: We show in this paper that a 3D scene can be revealed.

TITLE: Topology Reconstruction of Tree-Like Structure in Images via Structural Similarity Measure and Dominant Set Clustering

http://openaccess.thecvf.com/content_CVPR_2019/html/Xie_Topology_Reconstruction_of_Tree-Like_Structure_in_Images_via_Structural_Similarity_CVPR_2019_paper.html

AUTHORS: Jianyang Xie, Yitian Zhao, Yonghuai Liu, Pan Su, Yifan Zhao, Jun Cheng, Yalin Zheng, Jiang Liu

HIGHLIGHT: In this paper, we propose a novel curvilinear structural similarity measure to guide a dominant-set clustering approach to address this indispensable issue.

TITLE: Pyramidal Person Re-Identification via Multi-Loss Dynamic Training

http://openaccess.thecvf.com/content_CVPR_2019/html/Zheng_Pyramidal_Person_Re-Identification_via_Multi-Loss_Dynamic_Training_CVPR_2019_paper.html

AUTHORS: Feng Zheng, Cheng Deng, Xing Sun, Xinyang Jiang, Xiaowei Guo, Zongqiao Yu, Feiyue Huang, Rongrong Ji

HIGHLIGHT: Experimental results clearly demonstrate that the proposed method achieves the state-of-the-art results on three datasets.

TITLE: Holistic and Comprehensive Annotation of Clinically Significant Findings on Diverse CT Images: Learning From Radiology Reports and Label Ontology

http://openaccess.thecvf.com/content_CVPR_2019/html/Yan_Holistic_and_Comprehensive_Annotation_of_Clinically_Significant_Findings_on_Diverse_CT_Images_Learning_From_Radiology_Reports_and_Label_Ontology_CVPR_2019_paper.html

AUTHORS: Ke Yan, Yifan Peng, Veit Sandfort, Mohammadhadi Bagheri, Zhiyong Lu, Ronald M. Summers

HIGHLIGHT: In radiologists' routine work, one major task is to read a medical image, e.g., a CT scan, find significant lesions, and describe them in the radiology report.

TITLE: Robust Histopathology Image Analysis: To Label or to Synthesize?

http://openaccess.thecvf.com/content_CVPR_2019/html/Hou_Robust_Histopathology_Image_Analysis_To_Label_or_to_Synthesize_CVPR_2019_paper.html

AUTHORS: Le Hou, Ayush Agarwal, Dimitris Samaras, Tahsin M. Kurc, Rajarsi R. Gupta, Joel H. Saltz

HIGHLIGHT: We propose an unsupervised approach for histopathology image segmentation that synthesizes heterogeneous sets of training image patches, of every tissue type.

TITLE: Data Augmentation Using Learned Transformations for One-Shot Medical Image Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Data_Augmentation_Using_Learned_Transformations_for_One-Shot_Medical_Image_Segmentation_CVPR_2019_paper.html

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

HIGHLIGHT: We present an automated data augmentation method for synthesizing labeled medical images.

TITLE: Shifting More Attention to Video Salient Object Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Fan_Shifting_More_Attention_to_Video_Salient_Object_Detection_CVPR_2019_paper.html
AUTHORS: Deng-Ping Fan, Wenguan Wang, Ming-Ming Cheng, Jianbing Shen
HIGHLIGHT: With corresponding real human eye-fixation data, we obtain precise ground-truths.

TITLE: Neural Task Graphs: Generalizing to Unseen Tasks From a Single Video Demonstration
http://openaccess.thecvf.com/content_CVPR_2019/html/Huang_Neural_Task_Graphs_Generalizing_to_Unseen_Tasks_From_a_Single_CVPR_2019_paper.html
AUTHORS: De-An Huang, Suraj Nair, Danfei Xu, Yuke Zhu, Animesh Garg, Li Fei-Fei, Silvio Savarese, Juan Carlos Niebles
HIGHLIGHT: Our goal is to generate a policy to complete an unseen task given just a single video demonstration of the task in a given domain.

TITLE: Beyond Tracking: Selecting Memory and Refining Poses for Deep Visual Odometry
http://openaccess.thecvf.com/content_CVPR_2019/html/Xue_Beyond_Tracking_Selecting_Memory_and_Refining_Poses_for_Deep_Visual_CVPR_2019_paper.html
AUTHORS: Fei Xue, Xin Wang, Shunkai Li, Qiuyuan Wang, Junqiu Wang, Hongbin Zha
HIGHLIGHT: In contrast, we present a VO framework by incorporating two additional components called Memory and Refining.

TITLE: Image Generation From Layout
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Image_Generation_From_Layout_CVPR_2019_paper.html
AUTHORS: Bo Zhao, Lili Meng, Weidong Yin, Leonid Sigal
HIGHLIGHT: To address these challenges, we propose a novel approach for layout-based image generation; we call it Layout2Im.

TITLE: Multimodal Explanations by Predicting Counterfactuality in Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Kanehira_Multimodal_Explanations_by_Predicting_Counterfactuality_in_Videos_CVPR_2019_paper.html
AUTHORS: Atsushi Kanehira, Kentaro Takemoto, Sho Inayoshi, Tatsuya Harada
HIGHLIGHT: Requirements that the expected output should satisfy are referred to as counterfactuality in this paper: (1) Compatibility of visual-linguistic explanations, and (2) Positiveness/negativeness for the specific positive/negative class.

TITLE: Learning to Explain With Complemental Examples
http://openaccess.thecvf.com/content_CVPR_2019/html/Kanehira_Learning_to_Explain_With_Complemental_Examples_CVPR_2019_paper.html
AUTHORS: Atsushi Kanehira, Tatsuya Harada
HIGHLIGHT: We propose a novel framework to generate complemental explanations, on which the joint distribution of the variables to explain, and those to be explained is parameterized by three different neural networks: predictor, linguistic explainer, and example selector.

TITLE: HAQ: Hardware-Aware Automated Quantization With Mixed Precision
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_HAQ_Hardware-Aware_Automated_Quantization_With_Mixed_Precision_CVPR_2019_paper.html
AUTHORS: Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, Song Han
HIGHLIGHT: In this paper, we introduce the Hardware-Aware Automated Quantization (HAQ) framework which leverages the reinforcement learning to automatically determine the quantization policy, and we take the hardware accelerator's feedback in the design loop.

TITLE: Content Authentication for Neural Imaging Pipelines: End-To-End Optimization of Photo Provenance in Complex Distribution Channels
http://openaccess.thecvf.com/content_CVPR_2019/html/Korus_Content_Authentication_for_Neural_Imaging_Pipelines_End-To-End_Optimization_of_Photo_CVPR_2019_paper.html
AUTHORS: Pawel Korus, Nasir Memon
HIGHLIGHT: This paper explores end-to-end optimization of the entire image acquisition and distribution workflow to facilitate reliable forensic analysis at the end of the distribution channel.

TITLE: Inverse Procedural Modeling of Knitwear
http://openaccess.thecvf.com/content_CVPR_2019/html/Trunz_Inverse_Procedural_Modeling_of_Knitwear_CVPR_2019_paper.html
AUTHORS: Elena Trunz, Sebastian Merzbach, Jonathan Klein, Thomas Schulze, Michael Weinmann, Reinhard Klein

HIGHLIGHT: Inverse Procedural Modeling of Knitwear

TITLE: Estimating 3D Motion and Forces of Person-Object Interactions From Monocular Video
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Estimating_3D_Motion_and_Forces_of_Person-Object_Interactions_From_Monocular_CVPR_2019_paper.html

AUTHORS: Zongmian Li, Jiri Sedlar, Justin Carpentier, Ivan Laptev, Nicolas Mansard, Josef Sivic

HIGHLIGHT: In this paper, we introduce a method to automatically reconstruct the 3D motion of a person interacting with an object from a single RGB video.

TITLE: DeepMapping: Unsupervised Map Estimation From Multiple Point Clouds

http://openaccess.thecvf.com/content_CVPR_2019/html/Ding_DeepMapping_Unsupervised_Map_Estimation_From_Multiple_Point_Clouds_CVPR_2019_paper.html

AUTHORS: Li Ding, Chen Feng

HIGHLIGHT: We propose DeepMapping, a novel registration framework using deep neural networks (DNNs) as auxiliary functions to align multiple point clouds from scratch to a globally consistent frame.

TITLE: End-To-End Interpretable Neural Motion Planner

http://openaccess.thecvf.com/content_CVPR_2019/html/Zeng_End-To-End_Interpretable_Neural_Motion_Planner_CVPR_2019_paper.html

AUTHORS: Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, Raquel Urtasun

HIGHLIGHT: In this paper, we propose a neural motion planner for learning to drive autonomously in complex urban scenarios that include traffic-light handling, yielding, and interactions with multiple road-users.

TITLE: DuDoNet: Dual Domain Network for CT Metal Artifact Reduction

http://openaccess.thecvf.com/content_CVPR_2019/html/Lin_DuDoNet_Dual_Domain_Network_for_CT_Metal_Artifact_Reduction_CVPR_2019_paper.html

AUTHORS: Wei-An Lin, Haofu Liao, Cheng Peng, Xiaohang Sun, Jingdan Zhang, Jiebo Luo, Rama Chellappa, Shaohua Kevin Zhou

HIGHLIGHT: To overcome these difficulties, we propose an end-to-end trainable Dual Domain Network (DuDoNet) to simultaneously restore sinogram consistency and enhance CT images.

TITLE: Fast Spatio-Temporal Residual Network for Video Super-Resolution

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Fast_Spatio-Temporal_Residual_Network_for_Video_Super-Resolution_CVPR_2019_paper.html

AUTHORS: Sheng Li, Fengxiang He, Bo Du, Lefei Zhang, Yonghao Xu, Dacheng Tao

HIGHLIGHT: In this paper, we present a novel fast spatio-temporal residual network (FSTRN) to adopt 3D convolutions for the video SR task in order to enhance the performance while maintaining a low computational load.

TITLE: Complete the Look: Scene-Based Complementary Product Recommendation

http://openaccess.thecvf.com/content_CVPR_2019/html/Kang_Complete_the_Look_Scene-Based_Complementary_Product_Recommendation_CVPR_2019_paper.html

AUTHORS: Wang-Cheng Kang, Eric Kim, Jure Leskovec, Charles Rosenberg, Julian McAuley

HIGHLIGHT: In this work, we propose a new task called 'Complete the Look', which seeks to recommend visually compatible products based on scene images.

TITLE: Selective Sensor Fusion for Neural Visual-Inertial Odometry

http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Selective_Sensor_Fusion_for_Neural_Visual-Inertial_Odometry_CVPR_2019_paper.html

AUTHORS: Changhao Chen, Stefano Rosa, Yishu Miao, Chris Xiaoxuan Lu, Wei Wu, Andrew Markham, Niki Trigoni

HIGHLIGHT: We propose a novel end-to-end selective sensor fusion framework for monocular VIO, which fuses monocular images and inertial measurements in order to estimate the trajectory whilst improving robustness to real-life issues, such as missing and corrupted data or bad sensor synchronization.

TITLE: Look More Than Once: An Accurate Detector for Text of Arbitrary Shapes

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Look_More_Than_Once_An_Accurate_Detector_for_Text_of_CVPR_2019_paper.html

AUTHORS: Chengquan Zhang, Borong Liang, Zuming Huang, Mengyi En, Junyu Han, Errui Ding, Xinghao Ding

HIGHLIGHT: To address these two problems, we present a novel text detector namely LOMO, which localizes the text progressively for multiple times (or in other word, LOOK More than Once).

TITLE: Learning Binary Code for Personalized Fashion Recommendation

http://openaccess.thecvf.com/content_CVPR_2019/html/Lu_Learning_Binary_Code_for_Personalized_Fashion_Recommendation_CVPR_2019_paper.html

AUTHORS: Zhi Lu, Yang Hu, Yunchao Jiang, Yan Chen, Bing Zeng

HIGHLIGHT: In this paper, we propose to learn binary code for efficient personalized fashion outfits recommendation.

TITLE: Attention Based Glaucoma Detection: A Large-Scale Database and CNN Model

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Attention_Based_Glaucoma_Detection_A_Large-Scale_Database_and_CNN_Model_CVPR_2019_paper.html

AUTHORS: Liu Li, Mai Xu, Xiaofei Wang, Lai Jiang, Hanruo Liu

HIGHLIGHT: This paper proposes an attention-based CNN for glaucoma detection (AG-CNN).

TITLE: Privacy Protection in Street-View Panoramas Using Depth and Multi-View Imagery

http://openaccess.thecvf.com/content_CVPR_2019/html/Uittenbogaard_Privacy_Protection_in_Street-View_Panoramas_Using_Depth_and_Multi-View_Imagery_CVPR_2019_paper.html

AUTHORS: Ries Uittenbogaard, Clint Sebastian, Julien Vijverberg, Bas Boom, Dariu M. Gavrilă, Peter H.N. de With

HIGHLIGHT: In this paper, we propose a framework that is an alternative to blurring, which automatically removes and inpaints moving objects (e.g. pedestrians, vehicles) in street-view imagery.

TITLE: Grounding Human-To-Vehicle Advice for Self-Driving Vehicles

http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Grounding_Human-To-Vehicle_Advice_for_Self-Driving_Vehicles_CVPR_2019_paper.html

AUTHORS: Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, John Canny

HIGHLIGHT: Here, we propose to address this issue by augmenting training data with natural language advice from a human.

TITLE: Multi-Step Prediction of Occupancy Grid Maps With Recurrent Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Mohajerin_Multi-Step_Prediction_of_Occupancy_Grid_Maps_With_Recurrent_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Nima Mohajerin, Mohsen Rohani

HIGHLIGHT: We investigate the multi-step prediction of the drivable space, represented by Occupancy Grid Maps (OGMs), for autonomous vehicles.

TITLE: Connecting Touch and Vision via Cross-Modal Prediction

http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Connecting_Touch_and_Vision_via_Cross-Modal_Prediction_CVPR_2019_paper.html

AUTHORS: Yunzhu Li, Jun-Yan Zhu, Russ Tedrake, Antonio Torralba

HIGHLIGHT: In this work, we investigate the cross-modal connection between vision and touch.

TITLE: X2CT-GAN: Reconstructing CT From Biplanar X-Rays With Generative Adversarial Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Ying_X2CT-GAN_Reconstructing_CT_From_Biplanar_X-Rays_With_Generative_Adversarial_Networks_CVPR_2019_paper.html

AUTHORS: Xingde Ying, Heng Guo, Kai Ma, Jian Wu, Zhengxin Weng, Yefeng Zheng

HIGHLIGHT: In this work, we propose to reconstruct CT from two orthogonal X-rays using the generative adversarial network (GAN) framework.

TITLE: End-To-End Learned Random Walker for Seeded Image Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Cerrone_End-To-End_Learned_Random_Walker_for_Seeded_Image_Segmentation_CVPR_2019_paper.html

AUTHORS: Lorenzo Cerrone, Alexander Zeilmann, Fred A. Hamprecht

HIGHLIGHT: We present an end-to-end learned algorithm for seeded segmentation.

TITLE: Efficient Neural Network Compression

http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Efficient_Neural_Network_Compression_CVPR_2019_paper.html

AUTHORS: Hyeji Kim, Muhammad Umar Karim Khan, Chong-Min Kyung

HIGHLIGHT: In this paper we propose an efficient method for obtaining the rank configuration of the whole network.

TITLE: Cascaded Generative and Discriminative Learning for Microcalcification Detection in Breast Mammograms

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Cascaded_Generative_and_Discriminative_Learning_for_Microcalcification_Detection_in_Breast_CVPR_2019_paper.html

AUTHORS: Fandong Zhang, Ling Luo, Xinwei Sun, Zhen Zhou, Xiuli Li, Yizhou Yu, Yizhou Wang

HIGHLIGHT: In this paper, we propose a hybrid approach by taking advantages of both generative and discriminative models.

TITLE: C3AE: Exploring the Limits of Compact Model for Age Estimation
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_C3AE_Exploring_the_Limits_of_Compact_Model_for_Age_Estimation_CVPR_2019_paper.html
AUTHORS: Chao Zhang, Shuaicheng Liu, Xun Xu, Ce Zhu
HIGHLIGHT: In this work, we investigate the limits of compact model for small-scale image and propose an extremely Compact yet efficient Cascade Context-based Age Estimation model(C3AE).

TITLE: Adaptive Weighting Multi-Field-Of-View CNN for Semantic Segmentation in Pathology
http://openaccess.thecvf.com/content_CVPR_2019/html/Tokunaga_Adaptive_Weighting_Multi-Field-Of-View_CNN_for_Semantic_Segmentation_in_Pathology_CVPR_2019_paper.html
AUTHORS: Hiroki Tokunaga, Yuki Teramoto, Akihiko Yoshizawa, Ryoma Bise
HIGHLIGHT: In this paper, we propose a novel semantic segmentation method, called Adaptive-Weighting-Multi-Field-of-View-CNN (AWMF-CNN), that can adaptively use image features from images with different magnifications to segment multiple cancer subtype regions in the input image.

TITLE: In Defense of Pre-Trained ImageNet Architectures for Real-Time Semantic Segmentation of Road-Driving Images
http://openaccess.thecvf.com/content_CVPR_2019/html/Orsic_In_Defense_of_Pre-Trained_ImageNet_Architectures_for_Real-Time_Semantic_Segmentation_CVPR_2019_paper.html
AUTHORS: Marin Orsic, Ivan Kreso, Petra Bevanđić, Sinisa Segvić
HIGHLIGHT: We propose an alternative approach which achieves a significantly better performance across a wide range of computing budgets.

TITLE: Context-Aware Visual Compatibility Prediction
http://openaccess.thecvf.com/content_CVPR_2019/html/Cucurull_Context-Aware_Visual_Compatibility_Prediction_CVPR_2019_paper.html
AUTHORS: Guillem Cucurull, Perouz Taslakian, David Vazquez
HIGHLIGHT: In this work we propose a method that predicts compatibility between two items based on their visual features, as well as their context.

TITLE: Sim-To-Real via Sim-To-Sim: Data-Efficient Robotic Grasping via Randomized-To-Canonical Adaptation Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/James_Sim-To-Real_via_Sim-To-Sim_Data-Efficient_Robotic_Grasping_via_Randomized-To-Canonical_Adaptation_Networks_CVPR_2019_paper.html
AUTHORS: Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irpan, Julian Ibarz, Sergey Levine, Raia Hadsell, Konstantinos Bousmalis
HIGHLIGHT: In this paper, we present Randomized-to-Canonical Adaptation Networks (RCANs), a novel approach to crossing the visual reality gap that uses no real-world data.

TITLE: Multiview 2D/3D Rigid Registration via a Point-Of-Interest Network for Tracking and Triangulation
http://openaccess.thecvf.com/content_CVPR_2019/html/Liao_Multiview_2D3D_Rigid_Registration_via_a_Point-Of-Interest_Network_for_Tracking_CVPR_2019_paper.html
AUTHORS: HaoFu Liao, Wei-An Lin, Jiarui Zhang, Jingdan Zhang, Jiebo Luo, S. Kevin Zhou
HIGHLIGHT: We propose to tackle the problem of multiview 2D/3D rigid registration for intervention via a Point-Of-Interest Network for Tracking and Triangulation (POINT²).

TITLE: Context-Aware Spatio-Recurrent Curvilinear Structure Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Context-Aware_Spatio-Recurrent_Curvilinear_Structure_Segmentation_CVPR_2019_paper.html
AUTHORS: Feigege Wang, Yue Gu, Wenxi Liu, Yuanlong Yu, Shengfeng He, Jia Pan
HIGHLIGHT: In this paper, we propose a novel curvilinear structure segmentation approach using context-aware spatio-recurrent networks.

TITLE: An Alternative Deep Feature Approach to Line Level Keyword Spotting
http://openaccess.thecvf.com/content_CVPR_2019/html/Retsinas_An_Alternative_Deep_Feature_Approach_to_Line_Level_Keyword_Spotting_CVPR_2019_paper.html
AUTHORS: George Retsinas, Georgios Louloudis, Nikolaos Stamatopoulos, Giorgos Sfikas, Basilis Gatos
HIGHLIGHT: In this work, we propose a time and storage-efficient, deep feature-based approach that enables both the image and textual search options.

TITLE: Dynamics Are Important for the Recognition of Equine Pain in Video
http://openaccess.thecvf.com/content_CVPR_2019/html/Broome_Dynamics_Are_Important_for_the_Recognition_of_Equine_Pain_in_CVPR_2019_paper.html

AUTHORS: Sofia Broome, Karina Bech Gleerup, Pia Haubro Andersen, Hedvig Kjellstrom
HIGHLIGHT: In this study, we propose a deep recurrent two-stream architecture for the task of distinguishing pain from non-pain in videos of horses.

TITLE: LaserNet: An Efficient Probabilistic 3D Object Detector for Autonomous Driving
http://openaccess.thecvf.com/content_CVPR_2019/html/Meyer_LaserNet_An_Efficient_Probabilistic_3D_Object_Detector_for_Autonomous_Driving_CVPR_2019_paper.html
AUTHORS: Gregory P. Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, Carl K. Wellington
HIGHLIGHT: In this paper, we present LaserNet, a computationally efficient method for 3D object detection from LiDAR data for autonomous driving.

TITLE: Machine Vision Guided 3D Medical Image Compression for Efficient Transmission and Accurate Segmentation in the Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Machine_Vision_Guided_3D_Medical_Image_Compression_for_Efficient_Transmission_CVPR_2019_paper.html
AUTHORS: Zihao Liu, Xiaowei Xu, Tao Liu, Qi Liu, Yanzhi Wang, Yiyu Shi, Wujie Wen, Meiping Huang, Haiyun Yuan, Jian Zhuang
HIGHLIGHT: In this paper, we will use deep learning based medical image segmentation as a vehicle and demonstrate that interestingly, machine and human view the compression quality differently.

TITLE: PointPillars: Fast Encoders for Object Detection From Point Clouds
http://openaccess.thecvf.com/content_CVPR_2019/html/Lang_PointPillars_Fast_Encoders_for_Object_Detection_From_Point_Clouds_CVPR_2019_paper.html
AUTHORS: Alex H. Lang, Sourabh Vora, Holger Caesar, Luning Zhou, Jiong Yang, Oscar Beijbom
HIGHLIGHT: In this paper, we consider the problem of encoding a point cloud into a format appropriate for a downstream detection pipeline.

TITLE: Motion Estimation of Non-Holonomic Ground Vehicles From a Single Feature Correspondence Measured Over N Views
http://openaccess.thecvf.com/content_CVPR_2019/html/Huang_Motion_Estimation_of_Non-Holonomic_Ground_Vehicles_From_a_Single_Feature_CVPR_2019_paper.html
AUTHORS: Kun Huang, Yifu Wang, Laurent Kneip
HIGHLIGHT: We present the complete theory of this novel solver, and test it on both simulated and real data.

TITLE: From Coarse to Fine: Robust Hierarchical Localization at Large Scale
http://openaccess.thecvf.com/content_CVPR_2019/html/Sarlin_From_Coarse_to_Fine_Robust_Hierarchical_Localization_at_Large_Scale_CVPR_2019_paper.html
AUTHORS: Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, Marcin Dymczyk
HIGHLIGHT: In this paper we propose HF-Net, a hierarchical localization approach based on a monolithic CNN that simultaneously predicts local features and global descriptors for accurate 6-DoF localization.

TITLE: Large Scale High-Resolution Land Cover Mapping With Multi-Resolution Data
http://openaccess.thecvf.com/content_CVPR_2019/html/Robinson_Large_Scale_High-Resolution_Land_Cover_Mapping_With_Multi-Resolution_Data_CVPR_2019_paper.html
AUTHORS: Caleb Robinson, Le Hou, Kolya Malkin, Rachel Soobitsky, Jacob Czawlytko, Bistra Dilkina, Nebojsa Jojic
HIGHLIGHT: In this paper we propose multi-resolution data fusion methods for deep learning-based high-resolution land cover mapping from aerial imagery.

TITLE: Leveraging Heterogeneous Auxiliary Tasks to Assist Crowd Counting
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Leveraging_Heterogeneous_Auxiliary_Tasks_to_Assist_Crowd_Counting_CVPR_2019_paper.html
AUTHORS: Muming Zhao, Jian Zhang, Chongyang Zhang, Wenjun Zhang
HIGHLIGHT: In this paper, we propose to address these issues by leveraging the heterogeneous attributes compounded in the density map.