

TITLE: Unsupervised Deep Tracking
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Unsupervised_Deep_Tracking_CVPR_2019_paper.html
AUTHORS: Ning Wang, Yibing Song, Chao Ma, Wengang Zhou, Wei Liu, Houqiang Li
HIGHLIGHT: We propose an unsupervised visual tracking method in this paper.

TITLE: Tracking by Animation: Unsupervised Learning of Multi-Object Attentive Trackers
http://openaccess.thecvf.com/content_CVPR_2019/html/He_Tracking_by_Animation_Unsupervised_Learning_of_Multi-Object_Attentive_Trackers_CVPR_2019_paper.html
AUTHORS: Zhen He, Jian Li, Daxue Liu, Hangen He, David Barber
HIGHLIGHT: To achieve both label-free and end-to-end learning of MOT, we propose a Tracking-by-Animation framework, where a differentiable neural model first tracks objects from input frames and then animates these objects into reconstructed frames.

TITLE: Fast Online Object Tracking and Segmentation: A Unifying Approach
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Fast_Online_Object_Tracking_and_Segmentation_A_Unifying_Approach_CVPR_2019_paper.html
AUTHORS: Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, Philip H.S. Torr
HIGHLIGHT: In this paper we illustrate how to perform both visual object tracking and semi-supervised video object segmentation, in real-time, with a single simple approach.

TITLE: Object Tracking by Reconstruction With View-Specific Discriminative Correlation Filters
http://openaccess.thecvf.com/content_CVPR_2019/html/Kart_Object_Tracking_by_Reconstruction_With_View-Specific_Discriminative_Correlation_Filters_CVPR_2019_paper.html
AUTHORS: Ugur Kart, Alan Lukezic, Matej Kristan, Joni-Kristian Kamarainen, Jiri Matas
HIGHLIGHT: Object Tracking by Reconstruction With View-Specific Discriminative Correlation Filters

TITLE: SoPhie: An Attentive GAN for Predicting Paths Compliant to Social and Physical Constraints
http://openaccess.thecvf.com/content_CVPR_2019/html/Sadeghian_SoPhie_An_Attentive_GAN_for_Predicting_Paths_Compliant_to_Social_CVPR_2019_paper.html
AUTHORS: Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofighi, Silvio Savarese
HIGHLIGHT: We present SoPhie; an interpretable framework based on Generative Adversarial Network (GAN), which leverages two sources of information, the path history of all the agents in a scene, and the scene context information, using images of the scene.

TITLE: Leveraging Shape Completion for 3D Siamese Tracking
http://openaccess.thecvf.com/content_CVPR_2019/html/Giancola_Leveraging_Shape_Completion_for_3D_Siamese_Tracking_CVPR_2019_paper.html
AUTHORS: Silvio Giancola, Jesus Zarzar, Bernard Ghanem
HIGHLIGHT: In this paper, we investigate the versatility of Shape Completion for 3D Object Tracking in LIDAR point clouds.

TITLE: Target-Aware Deep Tracking
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Target-Aware_Deep_Tracking_CVPR_2019_paper.html
AUTHORS: Xin Li, Chao Ma, Baoyuan Wu, Zhenyu He, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a novel scheme to learn target-aware features, which can better recognize the targets undergoing significant appearance variations than pre-trained deep features.

TITLE: Spatiotemporal CNN for Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_Spatiotemporal_CNN_for_Video_Object_Segmentation_CVPR_2019_paper.html
AUTHORS: Kai Xu, Longyin Wen, Guorong Li, Liefeng Bo, Qingming Huang
HIGHLIGHT: In this paper, we present a unified, end-to-end trainable spatiotemporal CNN model for VOS, which consists of two branches, i.e., the temporal coherence branch and the spatial segmentation branch.

TITLE: Towards Rich Feature Discovery With Class Activation Maps Augmentation for Person Re-Identification
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Towards_Rich_Feature_Discovery_With_Class_Activation_Maps_Augmentation_for_CVPR_2019_paper.html
AUTHORS: Wenjie Yang, Houjing Huang, Zhang Zhang, Xiaotang Chen, Kaiqi Huang, Shu Zhang
HIGHLIGHT: This paper proposes to discover diverse discriminative visual cues without extra assistance, e.g., pose estimation, human parsing.

TITLE: Deep Rigid Instance Scene Flow

http://openaccess.thecvf.com/content_CVPR_2019/html/Ma_Deep_Rigid_Instance_Scene_Flow_CVPR_2019_paper.html

AUTHORS: Wei-Chiu Ma, Shenlong Wang, Rui Hu, Yuwen Xiong, Raquel Urtasun

HIGHLIGHT: In this paper we tackle the problem of scene flow estimation in the context of self-driving.

TITLE: See More, Know More: Unsupervised Video Object Segmentation With Co-Attention Siamese Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Lu_See_More_Know_More_Unsupervised_Video_Object_Segmentation_With_Co-Attention_CVPR_2019_paper.html

AUTHORS: Xiankai Lu, Wenguan Wang, Chao Ma, Jianbing Shen, Ling Shao, Fatih Porikli

HIGHLIGHT: We propose a unified and end-to-end trainable framework where different co-attention variants can be derived for mining the rich context within videos.

TITLE: Patch-Based Discriminative Feature Learning for Unsupervised Person Re-Identification

http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_Patch-Based_Discriminative_Feature_Learning_for_Unsupervised_Person_Re-Identification_CVPR_2019_paper.html

AUTHORS: Qize Yang, Hong-Xing Yu, Ancong Wu, Wei-Shi Zheng

HIGHLIGHT: In this work, we overcome this problem by proposing a patch-based unsupervised learning framework in order to learn discriminative feature from patches instead of the whole images.

TITLE: SPM-Tracker: Series-Parallel Matching for Real-Time Visual Object Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_SPM-Tracker_Series-Parallel_Matching_for_Real-Time_Visual_Object_Tracking_CVPR_2019_paper.html

AUTHORS: Guangting Wang, Chong Luo, Zhiwei Xiong, Wenjun Zeng

HIGHLIGHT: In this paper, we propose a SiamFC-based tracker, named SPM-Tracker, to tackle this challenge.

TITLE: SelfFlow: Self-Supervised Learning of Optical Flow

http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_SelfFlow_Self-Supervised_Learning_of_Optical_Flow_CVPR_2019_paper.html

AUTHORS: Pengpeng Liu, Michael Lyu, Irwin King, Jia Xu

HIGHLIGHT: We present a self-supervised learning approach for optical flow.

TITLE: Taking a Deeper Look at the Inverse Compositional Algorithm

http://openaccess.thecvf.com/content_CVPR_2019/html/Lv_Taking_a_Deeper_Look_at_the_Inverse_Compositional_Algorithm_CVPR_2019_paper.html

AUTHORS: Zhaoyang Lv, Frank Dellaert, James M. Rehg, Andreas Geiger

HIGHLIGHT: In this paper, we provide a modern synthesis of the classic inverse compositional algorithm for dense image alignment.

TITLE: Deeper and Wider Siamese Networks for Real-Time Visual Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Deep_and_Wider_Siamese_Networks_for_Real-Time_Visual_Tracking_CVPR_2019_paper.html

AUTHORS: Zhipeng Zhang, Houwen Peng

HIGHLIGHT: In this paper, we investigate how to leverage deeper and wider convolutional neural networks to enhance tracking robustness and accuracy.

TITLE: Self-Supervised Adaptation of High-Fidelity Face Models for Monocular Performance Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Yoon_Self-Supervised_Adaptation_of_High-Fidelity_Face_Models_for_Monocular_Performance_Tracking_CVPR_2019_paper.html

AUTHORS: Jae Shin Yoon, Takaaki Shiratori, Shoou-I Yu, Hyun Soo Park

HIGHLIGHT: In this paper, we propose a self-supervised domain adaptation approach to enable the animation of high-fidelity face models from a commodity camera.

TITLE: Diverse Generation for Multi-Agent Sports Games

http://openaccess.thecvf.com/content_CVPR_2019/html/Yeh_Diverse_Generation_for_Multi-Agent_Sports_Games_CVPR_2019_paper.html

AUTHORS: Raymond A. Yeh, Alexander G. Schwing, Jonathan Huang, Kevin Murphy

HIGHLIGHT: In this paper, we propose a new generative model for multi-agent trajectory data, focusing on the case of multi-player sports games.

TITLE: Efficient Online Multi-Person 2D Pose Tracking With Recurrent Spatio-Temporal Affinity Fields

http://openaccess.thecvf.com/content_CVPR_2019/html/Raaj_Efficient_Online_Multi-Person_2D_Pose_Tracking_With_Recurrent_Spatio-Temporal_Affinity_CVPR_2019_paper.html

AUTHORS: Yaadhav Raaj, Haroon Idrees, Gines Hidalgo, Yaser Sheikh
HIGHLIGHT: We present an online approach to efficiently and simultaneously detect and track 2D poses of multiple people in a video sequence.

TITLE: GFrames: Gradient-Based Local Reference Frame for 3D Shape Matching
http://openaccess.thecvf.com/content_CVPR_2019/html/Melzi_GFrames_Gradient-Based_Local_Reference_Frame_for_3D_Shape_Matching_CVPR_2019_paper.html
AUTHORS: Simone Melzi, Riccardo Spezialetti, Federico Tombari, Michael M. Bronstein, Luigi Di Stefano, Emanuele Rodola
HIGHLIGHT: We introduce GFrames, a novel local reference frame (LRF) construction for 3D meshes and point clouds.

TITLE: Eliminating Exposure Bias and Metric Mismatch in Multiple Object Tracking
http://openaccess.thecvf.com/content_CVPR_2019/html/Maksai_Eliminating_Exposure_Bias_and_Metric_Mismatch_in_Multiple_Object_Tracking_CVPR_2019_paper.html
AUTHORS: Andrii Maksai, Pascal Fua
HIGHLIGHT: In this paper, we introduce a new training procedure that confronts the algorithm to its own mistakes while explicitly attempting to minimize the number of switches, which results in better training.

TITLE: Graph Convolutional Tracking
http://openaccess.thecvf.com/content_CVPR_2019/html/Gao_Graph_Convolutional_Tracking_CVPR_2019_paper.html
AUTHORS: Junyu Gao, Tianzhu Zhang, Changsheng Xu
HIGHLIGHT: To comprehensively leverage the spatial-temporal structure of historical target exemplars and get benefit from the context information, in this work, we present a novel Graph Convolutional Tracking (GCT) method for high-performance visual tracking.

TITLE: ATOM: Accurate Tracking by Overlap Maximization
http://openaccess.thecvf.com/content_CVPR_2019/html/Danelljan_ATOM_Accurate_Tracking_by_Overlap_Maximization_CVPR_2019_paper.html
AUTHORS: Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, Michael Felsberg
HIGHLIGHT: We address this problem by proposing a novel tracking architecture, consisting of dedicated target estimation and classification components.

TITLE: Visual Tracking via Adaptive Spatially-Regularized Correlation Filters
http://openaccess.thecvf.com/content_CVPR_2019/html/Dai_Visual_Tracking_via_Adaptive_Spatially-Regularized_Correlation_Filters_CVPR_2019_paper.html
AUTHORS: Kenan Dai, Dong Wang, Huchuan Lu, Chong Sun, Jianhua Li
HIGHLIGHT: In this work, we propose a novel adaptive spatially-regularized correlation filters (ASRCF) model to simultaneously optimize the filter coefficients and the spatial regularization weight.

TITLE: Deep Tree Learning for Zero-Shot Face Anti-Spoofing
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Deep_Tree_Learning_for_Zero-Shot_Face_Anti-Spoofing_CVPR_2019_paper.html
AUTHORS: Yaojie Liu, Joel Stehouwer, Amin Jourabloo, Xiaoming Liu
HIGHLIGHT: In this work, we expand the ZSFA problem to a wide range of 13 types of spoof attacks, including print attack, replay attack, 3D mask attacks, and so on.

TITLE: ArcFace: Additive Angular Margin Loss for Deep Face Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.html
AUTHORS: Jiankang Deng, Jia Guo, Niannan Xue, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we propose an Additive Angular Margin Loss (ArcFace) to obtain highly discriminative features for face recognition.

TITLE: Learning Joint Gait Representation via Quintuplet Loss Minimization
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Learning_Joint_Gait_Representation_via_Quintuplet_Loss_Minimization_CVPR_2019_paper.html
AUTHORS: Kaihao Zhang, Wenhan Luo, Lin Ma, Wei Liu, Hongdong Li
HIGHLIGHT: In this paper, we propose a new Joint Unique-gait and Cross-gait Network (JUCNet), to combine the advantages of unique-gait representation with that of cross-gait representation, leading to an significantly improved performance.

TITLE: Gait Recognition via Disentangled Representation Learning

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Gait_Recognition_via_Disentangled_Representation_Learning_CVPR_2019_paper.html

AUTHORS: Ziyuan Zhang, Luan Tran, Xi Yin, Yousef Atoum, Xiaoming Liu, Jian Wan, Nanxin Wang
HIGHLIGHT: To remedy this issue, we propose a novel AutoEncoder framework to explicitly disentangle pose and appearance features from RGB imagery and the LSTM-based integration of pose features over time produces the gait feature.

TITLE: On the Continuity of Rotation Representations in Neural Networks

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhou_On_the_Continuity_of_Rotation_Representations_in_Neural_Networks_CVPR_2019_paper.html

AUTHORS: Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, Hao Li
HIGHLIGHT: In this paper, we advance a definition of a continuous representation, which can be helpful for training deep neural networks.

TITLE: Iterative Residual Refinement for Joint Optical Flow and Occlusion Estimation

http://openaccess.thecvf.com/content_CVPR_2019/html/Hur_Iterative_Residual_Refinement_for_Joint_Optical_Flow_and_Occlusion_Estimation_CVPR_2019_paper.html

AUTHORS: Junhwa Hur, Stefan Roth
HIGHLIGHT: Taking inspiration from both classical energy minimization approaches as well as residual networks, we propose an iterative residual refinement (IRR) scheme based on weight sharing that can be combined with several backbone networks.

TITLE: Inverse Discriminative Networks for Handwritten Signature Verification

http://openaccess.thecvf.com/content_CVPR_2019/html/Wei_Inverse_Discriminative_Networks_for_Handwritten_Signature_Verification_CVPR_2019_paper.html

AUTHORS: Ping Wei, Huan Li, Ping Hu
HIGHLIGHT: In this paper, we propose an inverse discriminative network (IDN) for writer-independent handwritten signature verification, which aims to determine whether a test signature is genuine or forged compared to the reference signature.

TITLE: Led3D: A Lightweight and Efficient Deep Approach to Recognizing Low-Quality 3D Faces

http://openaccess.thecvf.com/content_CVPR_2019/html/Mu_Led3D_A_Lightweight_and_Efficient_Deep_Approach_to_Recognizing_Low-Quality_CVPR_2019_paper.html

AUTHORS: Guodong Mu, Di Huang, Guosheng Hu, Jia Sun, Yunhong Wang
HIGHLIGHT: In this paper, we focus on 3D FR using low-quality data, targeting an efficient and accurate deep learning solution.

TITLE: ROI Pooled Correlation Filters for Visual Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_ROI_Pooled_Correlation_Filters_for_Visual_Tracking_CVPR_2019_paper.html

AUTHORS: Yuxuan Sun, Chong Sun, Dong Wang, You He, Huchuan Lu
HIGHLIGHT: In this paper, we propose a novel ROI pooled correlation filter (RPCF) algorithm for robust visual tracking.

TITLE: MOTS: Multi-Object Tracking and Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Voigtlaender_MOTS_Multi-Object_Tracking_and_Segmentation_CVPR_2019_paper.html

AUTHORS: Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, Bastian Leibe
HIGHLIGHT: We make our annotations, code, and models available at <https://www.vision.rwth-aachen.de/page/mots>.

TITLE: Siamese Cascaded Region Proposal Networks for Real-Time Visual Tracking

http://openaccess.thecvf.com/content_CVPR_2019/html/Fan_Siamese_Cascaded_Region_Proposal_Networks_for_Real-Time_Visual_Tracking_CVPR_2019_paper.html

AUTHORS: Heng Fan, Haibin Ling
HIGHLIGHT: Addressing these issues, we propose a multi-stage tracking framework, Siamese Cascaded RPN (C-RPN), which consists of a sequence of RPNs cascaded from deep high-level to shallow low-level layers in a Siamese network.

TITLE: PointFlowNet: Learning Representations for Rigid Motion Estimation From Point Clouds

http://openaccess.thecvf.com/content_CVPR_2019/html/Behl_PointFlowNet_Learning_Representations_for_Rigid_Motion_Estimation_From_Point_Clouds_CVPR_2019_paper.html

AUTHORS: Aseem Behl, Despoina Paschalidou, Simon Donne, Andreas Geiger
HIGHLIGHT: In this paper, we propose to estimate 3D motion from such unstructured point clouds using a deep neural network.

TITLE: Object Discovery in Videos as Foreground Motion Clustering
http://openaccess.thecvf.com/content_CVPR_2019/html/Xie_Object_Discovery_in_Videos_as_Foreground_Motion_Clustering_CVPR_2019_paper.html

AUTHORS: Christopher Xie, Yu Xiang, Zaid Harchaoui, Dieter Fox

HIGHLIGHT: We consider the problem of providing dense segmentation masks for object discovery in videos.

TITLE: Towards Natural and Accurate Future Motion Prediction of Humans and Animals
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Towards_Natural_and_Accurate_Future_Motion_Prediction_of_Humans_and_CVPR_2019_paper.html

AUTHORS: Zhenguang Liu, Shuang Wu, Shuyuan Jin, Qi Liu, Shijian Lu, Roger Zimmermann, Li Cheng

HIGHLIGHT: To address these problems, we propose to explicitly encode anatomical constraints by modeling their skeletons with a Lie algebra representation.

TITLE: Automatic Face Aging in Videos via Deep Reinforcement Learning
http://openaccess.thecvf.com/content_CVPR_2019/html/Duong_Automatic_Face_Aging_in_Videos_via_Deep_Reinforcement_Learning_CVPR_2019_paper.html

AUTHORS: Chi Nhan Duong, Khoa Luu, Kha Gia Quach, Nghia Nguyen, Eric Patterson, Tien D. Bui, Ngan Le

HIGHLIGHT: This paper presents a novel approach for synthesizing automatically age-progressed facial images in video sequences using Deep Reinforcement Learning.

TITLE: Multi-Adversarial Discriminative Deep Domain Generalization for Face Presentation Attack Detection

http://openaccess.thecvf.com/content_CVPR_2019/html/Shao_Multi-Adversarial_Discriminative_Deep_Domain_Generalization_for_Face_Presentation_Attack_Detection_CVPR_2019_paper.html

AUTHORS: Rui Shao, Xiangyuan Lan, Jiawei Li, Pong C. Yuen

HIGHLIGHT: We propose to learn a generalized feature space via a novel multi-adversarial discriminative deep domain generalization framework.

TITLE: SR-LSTM: State Refinement for LSTM Towards Pedestrian Trajectory Prediction

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_SR-LSTM_State_Refinement_for_LSTM_Towards_Pedestrian_Trajectory_Prediction_CVPR_2019_paper.html

AUTHORS: Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, Nanning Zheng

HIGHLIGHT: In order to address this issue, we propose a data-driven state refinement module for LSTM network (SR-LSTM), which activates the utilization of the current intention of neighbors, and jointly and iteratively refines the current states of all participants in the crowd through a message passing mechanism.

TITLE: Unsupervised Deep Epipolar Flow for Stationary or Dynamic Scenes

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhong_Unsupervised_Deep_Epipolar_Flow_for_Stationary_or_Dynamic_Scenes_CVPR_2019_paper.html

AUTHORS: Yiran Zhong, Pan Ji, Jianyuan Wang, Yuchao Dai, Hongdong Li

HIGHLIGHT: In this paper, we propose Deep Epipolar Flow, an unsupervised optical flow method which incorporates global geometric constraints into network learning.

TITLE: An Efficient Schmidt-EKF for 3D Visual-Inertial SLAM

http://openaccess.thecvf.com/content_CVPR_2019/html/Geneva_An_Efficient_Schmidt-EKF_for_3D_Visual-Inertial_SLAM_CVPR_2019_paper.html

AUTHORS: Patrick Geneva, James Maley, Guoquan Huang

HIGHLIGHT: In this paper, we propose a novel, high-precision, efficient visual-inertial (VI)-SLAM algorithm, termed Schmidt-EKF VI-SLAM (SEVIS), which optimally fuses IMU measurements and monocular images in a tightly-coupled manner to provide 3D motion tracking with bounded error.

TITLE: A Neural Temporal Model for Human Motion Prediction

http://openaccess.thecvf.com/content_CVPR_2019/html/Gopalakrishnan_A_Neural_Temporal_Model_for_Human_Motion_Prediction_CVPR_2019_paper.html

AUTHORS: Anand Gopalakrishnan, Ankur Mali, Dan Kifer, Lee Giles, Alexander G. Ororbia

HIGHLIGHT: We propose novel neural temporal models for predicting and synthesizing human motion, achieving state-of-the-art in modeling long-term motion trajectories while being competitive with prior work in short-term prediction and requiring significantly less computation.

TITLE: Multi-Agent Tensor Fusion for Contextual Trajectory Prediction

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Multi-Agent_Tensor_Fusion_for_Contextual_Trajectory_Prediction_CVPR_2019_paper.html

AUTHORS: Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, Ying Nian Wu

HIGHLIGHT: Experiments on both highway driving and pedestrian crowd datasets show that the model achieves state-of-the-art prediction accuracy.