

1, TITLE: Deep Compositional Captioning: Describing Novel Object Categories Without Paired Training Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Hendricks_Deep_Compositional_Captioning_CVPR_2016_paper.html
AUTHORS: Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond Mooney, Kate Saenko, Trevor Darrell
HIGHLIGHT: In this work, we propose the Deep Compositional Captioner (DCC) to address the task of generating descriptions of novel objects which are not present in paired image-sentence datasets.

2, TITLE: Generation and Comprehension of Unambiguous Object Descriptions
http://openaccess.thecvf.com/content_cvpr_2016/html/Mao_Generation_and_Comprehension_CVPR_2016_paper.html
AUTHORS: Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, Kevin Murphy
HIGHLIGHT: We propose a method that can generate an unambiguous description (known as a referring expression) of a specific object or region in an image, and which can also comprehend or interpret such an expression to infer which object is being described.
We also present a new large-scale dataset for referring expressions, based on MS-COCO.
We have released the dataset and a toolbox for visualization and evaluation, see https://github.com/mjhucla/Google_Refexp_toolbox.

3, TITLE: Stacked Attention Networks for Image Question Answering
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Stacked_Attention_Networks_CVPR_2016_paper.html
AUTHORS: Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Smola
HIGHLIGHT: This paper presents stacked attention networks (SANs) that learn to answer natural language questions from images.

4, TITLE: Image Question Answering Using Convolutional Neural Network With Dynamic Parameter Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Noh_Image_Question_Answering_CVPR_2016_paper.html
AUTHORS: Hyeonwoo Noh, Paul Hongsuck Seo, Bohyung Han
HIGHLIGHT: We tackle image question answering (ImageQA) problem by learning a convolutional neural network (CNN) with a dynamic parameter layer whose weights are determined adaptively based on questions.

5, TITLE: Neural Module Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Andreas_Neural_Module_Networks_CVPR_2016_paper.html
AUTHORS: Jacob Andreas, Marcus Rohrbach, Trevor Darrell, Dan Klein
HIGHLIGHT: We describe a procedure for constructing and learning neural module networks, which compose collections of jointly-trained neural "modules" into deep networks for question answering.

6, TITLE: Learning Deep Representations of Fine-Grained Visual Descriptions
http://openaccess.thecvf.com/content_cvpr_2016/html/Reed_Learning_Deep_Representations_CVPR_2016_paper.html
AUTHORS: Scott Reed, Zeynep Akata, Honglak Lee, Bernt Schiele
HIGHLIGHT: We propose to overcome these limitations by training neural language models from scratch; i.e. without pre-training and only consuming words and characters.

7, TITLE: Multi-Cue Zero-Shot Learning With Strong Supervision
http://openaccess.thecvf.com/content_cvpr_2016/html/Akata_Multi-Cue_Zero-Shot_Learning_CVPR_2016_paper.html
AUTHORS: Zeynep Akata, Mateusz Malinowski, Mario Fritz, Bernt Schiele
HIGHLIGHT: We achieve our goal by a joint embedding framework that maps multiple text parts as well as multiple semantic parts into a common space.

8, TITLE: Latent Embeddings for Zero-Shot Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Xian_Latent_Embeddings_for_CVPR_2016_paper.html
AUTHORS: Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, Bernt Schiele
HIGHLIGHT: We present a novel latent embedding model for learning a compatibility function between image and class embeddings, in the context of zero-shot classification.

9, TITLE: One-Shot Learning of Scene Locations via Feature Trajectory Transfer
http://openaccess.thecvf.com/content_cvpr_2016/html/Kwitt_One-Shot_Learning_of_CVPR_2016_paper.html
AUTHORS: Roland Kwitt, Sebastian Hegenbart, Marc Niethammer
HIGHLIGHT: In this work, we investigate the variability in these transient attributes as a rich source of information for studying how image representations change as a function of attribute strength.

10, TITLE: Learning Attributes Equals Multi-Source Domain Generalization
http://openaccess.thecvf.com/content_cvpr_2016/html/Gan_Learning_Attributes_Equals_CVPR_2016_paper.html
AUTHORS: Chuang Gan, Tianbao Yang, Boqing Gong

HIGHLIGHT: Noting that this is analogous to the objective of multi-source domain generalization, if we treat each category as a domain, we provide a novel perspective to attribute detection and propose to gear the techniques in multi-source domain generalization for the purpose of learning cross-category generalizable attribute detectors.

11, **TITLE:** Anticipating Visual Representations From Unlabeled Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Vondrick_Anticipating_Visual_Representations_CVPR_2016_paper.html
AUTHORS: Carl Vondrick, Hamed Pirsiavash, Antonio Torralba
HIGHLIGHT: We present a framework that capitalizes on temporal structure in unlabeled video to learn to anticipate human actions and objects.

12, **TITLE:** Learning to Assign Orientations to Feature Points
http://openaccess.thecvf.com/content_cvpr_2016/html/Yi_Learning_to_Assign_CVPR_2016_paper.html
AUTHORS: Kwang Moo Yi, Yannick Verdie, Pascal Fua, Vincent Lepetit
HIGHLIGHT: To avoid the tedious and almost impossible task of finding a target orientation to learn, we propose to use Siamese networks which implicitly find the optimal orientations during training.

13, **TITLE:** Learning Dense Correspondence via 3D-Guided Cycle Consistency
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Learning_Dense_Correspondence_CVPR_2016_paper.html
AUTHORS: Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, Alexei A. Efros
HIGHLIGHT: For each pair of training images we find an appropriate 3D CAD model and render two synthetic views to link in with the pair, establishing a correspondence flow 4-cycle.

14, **TITLE:** The Global Patch Collider
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_The_Global_Patch_CVPR_2016_paper.html
AUTHORS: Shenglong Wang, Sean Ryan Fanello, Christoph Rhemann, Shahram Izadi, Pushmeet Kohli
HIGHLIGHT: This paper proposes a novel extremely efficient, fully-parallelizable, task-specific algorithm for the computation of global point-wise correspondences in images and videos.

15, **TITLE:** Joint Probabilistic Matching Using m-Best Solutions
http://openaccess.thecvf.com/content_cvpr_2016/html/Rezatofighi_Joint_Probabilistic_Matching_CVPR_2016_paper.html
AUTHORS: Seyed Hamid Rezatofighi, Anton Milan, Zhen Zhang, Qinfeng Shi, Anthony Dick, Ian Reid
HIGHLIGHT: Here, we propose a generic approach to efficiently approximate the marginal distributions by exploiting the m-best solutions of the original problem.

16, **TITLE:** Face Alignment Across Large Poses: A 3D Solution
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Face_Alignment_Across_CVPR_2016_paper.html
AUTHORS: Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, Stan Z. Li
HIGHLIGHT: In this paper, we propose a solution to the three problems in a new alignment framework, called 3D Dense Face Alignment (3DDFA), in which a dense 3D face model is fitted to the image via convolutional neural network (CNN).

17, **TITLE:** Interactive Segmentation on RGBD Images via Cue Selection
http://openaccess.thecvf.com/content_cvpr_2016/html/Feng_Interactive_Segmentation_on_CVPR_2016_paper.html
AUTHORS: Jie Feng, Brian Price, Scott Cohen, Shih-Fu Chang
HIGHLIGHT: We propose a novel interactive segmentation algorithm which can incorporate multiple feature cues like color, depth, and normals in a unified graph cut framework to leverage these cues more effectively.

18, **TITLE:** Layered Scene Decomposition via the Occlusion-CRF
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Layered_Scene_Decomposition_CVPR_2016_paper.html
AUTHORS: Chen Liu, Pushmeet Kohli, Yasutaka Furukawa
HIGHLIGHT: We propose a novel "Occlusion-CRF" model that allows for the integration of sophisticated priors to regularize the solution space and enables the automatic inference of the layer decomposition.

19, **TITLE:** Affinity CNN: Learning Pixel-Centric Pairwise Relations for Figure/Ground Embedding
http://openaccess.thecvf.com/content_cvpr_2016/html/Maire_Affinity_CNN_Learning_CVPR_2016_paper.html
AUTHORS: Michael Maire, Takuya Narihira, Stella X. Yu
HIGHLIGHT: Spectral embedding provides a framework for solving perceptual organization problems, including image segmentation and figure/ground organization.

20, **TITLE:** Weakly Supervised Object Boundaries
http://openaccess.thecvf.com/content_cvpr_2016/html/Khoreva_Weakly_Supervised_Object_CVPR_2016_paper.html

AUTHORS: Anna Khoreva, Rodrigo Benenson, Mohamed Omran, Matthias Hein, Bernt Schiele
HIGHLIGHT: In this paper we propose a technique to generate weakly supervised annotations and show that bounding box annotations alone suffice to reach high-quality object boundaries without using any object-specific boundary annotations.

21, TITLE: Object Contour Detection With a Fully Convolutional Encoder-Decoder Network
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Object_Contour_Detection_CVPR_2016_paper.html
AUTHORS: Jimei Yang, Brian Price, Scott Cohen, Honglak Lee, Ming-Hsuan Yang
HIGHLIGHT: We develop a deep learning algorithm for contour detection with a fully convolutional encoder-decoder network.

22, TITLE: What Value Do Explicit High Level Concepts Have in Vision to Language Problems?
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_What_Value_Do_CVPR_2016_paper.html
AUTHORS: Qi Wu, Chunhua Shen, Lingqiao Liu, Anthony Dick, Anton van den Hengel
HIGHLIGHT: In this paper we investigate whether this direct approach succeeds due to, or despite, the fact that it avoids the explicit representation of high-level information.

23, TITLE: Fast Detection of Curved Edges at Low SNR
http://openaccess.thecvf.com/content_cvpr_2016/html/Ofir_Fast_Detection_of_CVPR_2016_paper.html
AUTHORS: Nati Ofir, Meirav Galun, Boaz Nadler, Ronen Basri
HIGHLIGHT: In this paper we develop a novel multiscale method to detect curved edges in noisy images.

24, TITLE: Object Skeleton Extraction in Natural Images by Fusing Scale-Associated Deep Side Outputs
http://openaccess.thecvf.com/content_cvpr_2016/html/Shen_Object_Skeleton_Extraction_CVPR_2016_paper.html
AUTHORS: Wei Shen, Kai Zhao, Yuan Jiang, Yan Wang, Zhijiang Zhang, Xiang Bai
HIGHLIGHT: In this paper, we present a fully convolutional network with multiple scale-associated side outputs to address this problem.

25, TITLE: Learning Relaxed Deep Supervision for Better Edge Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Learning_Relaxed_Deep_CVPR_2016_paper.html
AUTHORS: Yu Liu, Michael S. Lew
HIGHLIGHT: We propose using relaxed deep supervision (RDS) within convolutional neural networks for edge detection.

26, TITLE: Occlusion Boundary Detection via Deep Exploration of Context
http://openaccess.thecvf.com/content_cvpr_2016/html/Fu_Occlusion_Boundary_Detection_CVPR_2016_paper.html
AUTHORS: Huan Fu, Chaohui Wang, Dacheng Tao, Michael J. Black
HIGHLIGHT: In this paper, we improve occlusion boundary detection via enhanced exploration of contextual information (e.g., local structural boundary patterns, observations from surrounding regions, and temporal context), and in doing so develop a novel approach based on convolutional neural networks (CNNs) and conditional random fields (CRFs).

27, TITLE: SemiContour: A Semi-Supervised Learning Approach for Contour Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_SemiContour_A_Semi-Supervised_CVPR_2016_paper.html
AUTHORS: Zizhao Zhang, Fuyong Xing, Xiaoshuang Shi, Lin Yang
HIGHLIGHT: In this paper, we investigate the usage of semi-supervised learning (SSL) to obtain competitive detection accuracy with very limited training data (three labeled images).

28, TITLE: Learning to Localize Little Landmarks
http://openaccess.thecvf.com/content_cvpr_2016/html/Singh_Learning_to_Localize_CVPR_2016_paper.html
AUTHORS: Saurabh Singh, Derek Hoiem, David Forsyth
HIGHLIGHT: We describe a method to find such landmarks by finding a sequence of latent landmarks, each with a prediction model.

29, TITLE: InterActive: Inter-Layer Activeness Propagation
http://openaccess.thecvf.com/content_cvpr_2016/html/Xie_InterActive_Inter-Layer_Activeness_CVPR_2016_paper.html
AUTHORS: Lingxi Xie, Liang Zheng, Jingdong Wang, Alan L. Yuille, Qi Tian
HIGHLIGHT: In this paper, we present InterActive, a novel algorithm which computes the activeness of neurons and network connections.

30, TITLE: Exploit Bounding Box Annotations for Multi-Label Object Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Exploit_Bounding_Box_CVPR_2016_paper.html
AUTHORS: Hao Yang, Joey Tianyi Zhou, Yu Zhang, Bin-Bin Gao, Jianxin Wu, Jianfei Cai

HIGHLIGHT: In this paper, we incorporate local information to enhance the feature discriminative power.

31, **TITLE:** TI-Pooling: Transformation-Invariant Pooling for Feature Learning in Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Laptev_TI-Pooling_Transformation-Invariant_Pooling_CVPR_2016_paper.html

AUTHORS: Dmitry Laptev, Nikolay Savinov, Joachim M. Buhmann, Marc Pollefeys

HIGHLIGHT: In this paper we present a deep neural network topology that incorporates a simple to implement transformation-invariant pooling operator (TI-pooling).

32, **TITLE:** Fashion Style in 128 Floats: Joint Ranking and Classification Using Weak Data for Feature Extraction
http://openaccess.thecvf.com/content_cvpr_2016/html/Simo-Serra_Fashion_Style_in_CVPR_2016_paper.html

AUTHORS: Edgar Simo-Serra, Hiroshi Ishikawa

HIGHLIGHT: We propose a novel approach for learning features from weakly-supervised data by joint ranking and classification.

33, **TITLE:** Equiangular Kernel Dictionary Learning With Applications to Dynamic Texture Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Quan_Equiangular_Kernel_Dictionary_CVPR_2016_paper.html

AUTHORS: Yuhui Quan, Chenglong Bao, Hui Ji

HIGHLIGHT: This paper proposed an equiangular kernel dictionary learning method with optimal mutual coherence to exploit the nonlinear sparsity of high-dimensional visual data.

34, **TITLE:** Compact Bilinear Pooling

http://openaccess.thecvf.com/content_cvpr_2016/html/Gao_Compact_Bilinear_Pooling_CVPR_2016_paper.html

AUTHORS: Yang Gao, Oscar Beijbom, Ning Zhang, Trevor Darrell

HIGHLIGHT: We propose two compact bilinear representations with the same discriminative power as the full bilinear representation but with only a few thousand dimensions.

35, **TITLE:** Accumulated Stability Voting: A Robust Descriptor From Descriptors of Multiple Scales

http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Accumulated_Stability_Voting_CVPR_2016_paper.html

AUTHORS: Tsun-Yi Yang, Yen-Yu Lin, Yung-Yu Chuang

HIGHLIGHT: This paper proposes a novel local descriptor through accumulated stability voting (ASV).

36, **TITLE:** CoMaL: Good Features to Match on Object Boundaries

http://openaccess.thecvf.com/content_cvpr_2016/html/Ravindran_CoMaL_Good_Features_CVPR_2016_paper.html

AUTHORS: Swarna K. Ravindran, Anurag Mittal

HIGHLIGHT: In this paper, we propose a new approach for feature detection, tracking and re-detection that gives significantly improved results at the object boundaries.

37, **TITLE:** Progressive Feature Matching With Alternate Descriptor Selection and Correspondence Enrichment

http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Progressive_Feature_Matching_CVPR_2016_paper.html

AUTHORS: Yuan-Ting Hu, Yen-Yu Lin

HIGHLIGHT: We aim at tackling the two issues by integrating adaptive descriptor selection and progressive candidate enrichment into image matching.

38, **TITLE:** A New Finsler Minimal Path Model With Curvature Penalization for Image Segmentation and Closed Contour Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_A_New_Finsler_CVPR_2016_paper.html

AUTHORS: Da Chen, Jean-Marie Mirebeau, Laurent D. Cohen

HIGHLIGHT: In this paper, we propose a new curvature penalized minimal path model for image segmentation via closed contour detection based on the weighted Euler elastica curves, firstly introduced to the field of computer vision in [22].

39, **TITLE:** Scale-Aware Alignment of Hierarchical Image Segmentation

http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Scale-Aware_Alignment_of_CVPR_2016_paper.html

AUTHORS: Yuhua Chen, Dengxin Dai, Jordi Pont-Tuset, Luc Van Gool

HIGHLIGHT: Image segmentation is a key component in many computer vision systems, and it is recovering a prominent spot in the literature as methods improve and overcome their limitations.

40, **TITLE:** Deep Interactive Object Selection

http://openaccess.thecvf.com/content_cvpr_2016/html/Xu_Deep_Interactive_Object_CVPR_2016_paper.html

AUTHORS: Ning Xu, Brian Price, Scott Cohen, Jimei Yang, Thomas S. Huang

HIGHLIGHT: In this paper, we present a novel deep-learning-based algorithm which has much better understanding of objectness and can reduce user interactions to just a few clicks.

41, **TITLE:** Pull the Plug? Predicting If Computers or Humans Should Segment Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Gurari_Pull_the_Plug_CVPR_2016_paper.html

AUTHORS: Danna Gurari, Suyog Jain, Margrit Betke, Kristen Grauman

HIGHLIGHT: We propose a resource allocation framework for predicting how best to allocate a fixed budget of human annotation effort in order to collect higher quality segmentations for a given batch of images and automated methods.

42, **TITLE:** In the Shadows, Shape Priors Shine: Using Occlusion to Improve Multi-Region Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Kihara_In_the_Shadows_CVPR_2016_paper.html

AUTHORS: Yuka Kihara, Matvey Soloviev, Tshuan Chen

HIGHLIGHT: We present a new algorithm for multi-region segmentation of 2D images with objects that may partially occlude each other.

43, **TITLE:** Convexity Shape Constraints for Image Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Royer_Convexity_Shape_Constraints_CVPR_2016_paper.html

AUTHORS: Loic A. Royer, David L. Richmond, Carsten Rother, Bjoern Andres, Dagmar Kainmueller

HIGHLIGHT: In this work, we introduce a new approach that allows, for the first time, to constrain some or all components of a segmentation to have convex shapes.

44, **TITLE:** MCMC Shape Sampling for Image Segmentation With Nonparametric Shape Priors
http://openaccess.thecvf.com/content_cvpr_2016/html/Erdil_MCMC_Shape_Sampling_CVPR_2016_paper.html

AUTHORS: Ertunc Erdil, Sinan Yildirim, Mujdat Cetin, Tolga Tasdizen

HIGHLIGHT: For such characterization, we propose a Markov chain Monte Carlo (MCMC) sampling-based image segmentation algorithm that uses statistical shape priors.

45, **TITLE:** From Noise Modeling to Blind Image Denoising
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_From_Noise_Modeling_CVPR_2016_paper.html

AUTHORS: Fengyuan Zhu, Guangyong Chen, Pheng-Ann Heng

HIGHLIGHT: As the number of components for MoG is unknown practically, this work adopts Bayesian nonparametric technique and proposes a novel Low-rank MoG filter (LR-MoG) to recover clean signals (patches) from noisy ones contaminated by MoG noise.

46, **TITLE:** Efficient and Robust Color Consistency for Community Photo Collections
http://openaccess.thecvf.com/content_cvpr_2016/html/Park_Efficient_and_Robust_CVPR_2016_paper.html

AUTHORS: Jaesik Park, Yu-Wing Tai, Sudipta N. Sinha, In So Kweon

HIGHLIGHT: We present an efficient technique to optimize color consistency of a collection of images depicting a common scene.

47, **TITLE:** Needle-Match: Reliable Patch Matching Under High Uncertainty
http://openaccess.thecvf.com/content_cvpr_2016/html/Lotan_Needle-Match_Reliable_Patch_CVPR_2016_paper.html

AUTHORS: Or Lotan, Michal Irani

HIGHLIGHT: In this paper we present a patch representation called "Needle", which consists of small multi-scale versions of the patch and its immediate surrounding region.

48, **TITLE:** ReconNet: Non-Iterative Reconstruction of Images From Compressively Sensed Measurements
http://openaccess.thecvf.com/content_cvpr_2016/html/Kulkarni_ReconNet_Non-Iterative_Reconstruction_CVPR_2016_paper.html

AUTHORS: Kuldeep Kulkarni, Suhas Lohit, Pavan Turaga, Ronan Kerviche, Amit Ashok

HIGHLIGHT: The goal of this paper is to present a non-iterative and more importantly an extremely fast algorithm to reconstruct images from compressively sensed (CS) random measurements.

49, **TITLE:** Soft-Segmentation Guided Object Motion Deblurring
http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Soft-Segmentation_Guided_Object_CVPR_2016_paper.html

AUTHORS: Jinshan Pan, Zhe Hu, Zhixun Su, Hsin-Ying Lee, Ming-Hsuan Yang

HIGHLIGHT: To address these problems, we propose a novel model for object motion deblurring.

50, **TITLE:** Two Illuminant Estimation and User Correction Preference
http://openaccess.thecvf.com/content_cvpr_2016/html/Cheng_Two_Illuminant_Estimation_CVPR_2016_paper.html

AUTHORS: Dongliang Cheng, Abdelrahman Abdelhamed, Brian Price, Scott Cohen, Michael S. Brown

- HIGHLIGHT:** This paper examines the problem of white-balance correction when a scene contains two illuminations.
- 51, **TITLE:** Deep Contrast Learning for Salient Object Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Deep_Contrast_Learning_CVPR_2016_paper.html
AUTHORS: Guanbin Li, Yizhou Yu
HIGHLIGHT: In this paper, we propose an end-to-end deep contrast network to overcome the aforementioned limitations.
- 52, **TITLE:** Multiview Image Completion With Space Structure Propagation
http://openaccess.thecvf.com/content_cvpr_2016/html/Baek_Multiview_Image_Completion_CVPR_2016_paper.html
AUTHORS: Seung-Hwan Baek, Inchang Choi, Min H. Kim
HIGHLIGHT: We present a multiview image completion method that provides geometric consistency among different views by propagating space structures.
- 53, **TITLE:** Composition-Preserving Deep Photo Aesthetics Assessment
http://openaccess.thecvf.com/content_cvpr_2016/html/Mai_Composition-Preserving_Deep_Photo_CVPR_2016_paper.html
AUTHORS: Long Mai, Hailin Jin, Feng Liu
HIGHLIGHT: In this paper, we present a composition-preserving deep ConvNet method that directly learns aesthetics features from the original input images without any image transformations.
- 54, **TITLE:** Automatic Image Cropping : A Computational Complexity Study
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Automatic_Image_Cropping_CVPR_2016_paper.html
AUTHORS: Jiansheng Chen, Gaocheng Bai, Shaoheng Liang, Zhengqin Li
HIGHLIGHT: We demonstrate that under appropriate formulations, this task can be achieved using efficient algorithms with low computational complexity.
- 55, **TITLE:** A Deeper Look at Saliency: Feature Contrast, Semantics, and Beyond
http://openaccess.thecvf.com/content_cvpr_2016/html/Bruce_A_Deeper_Look_CVPR_2016_paper.html
AUTHORS: Neil D. B. Bruce, Christopher Catton, Sasa Janjic
HIGHLIGHT: In this paper we consider the problem of visual saliency modeling, including both human gaze prediction and salient object segmentation.
- 56, **TITLE:** Spatially Binned ROC: A Comprehensive Saliency Metric
http://openaccess.thecvf.com/content_cvpr_2016/html/Wloka_Spatially_Binned_ROC_CVPR_2016_paper.html
AUTHORS: Calden Wloka, John Tsotsos
HIGHLIGHT: To capture and quantify these known sources of bias, we propose a novel metric for measuring saliency algorithm performance: the spatially binned ROC (spROC).
- 57, **TITLE:** GraB: Visual Saliency via Novel Graph Model and Background Priors
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_GraB_Visual_Saliency_CVPR_2016_paper.html
AUTHORS: Qiaosong Wang, Wen Zheng, Robinson Piramuthu
HIGHLIGHT: We propose an unsupervised bottom-up saliency detection approach by exploiting novel graph structure and background priors.
- 58, **TITLE:** Predicting When Saliency Maps Are Accurate and Eye Fixations Consistent
http://openaccess.thecvf.com/content_cvpr_2016/html/Volokitin_Predicting_When_Saliency_CVPR_2016_paper.html
AUTHORS: Anna Volokitin, Michael Gygli, Xavier Boix
HIGHLIGHT: In this paper, we show that using features from DCNNs for object recognition we can make predictions that enrich the information provided by saliency models.
- 59, **TITLE:** Split and Match: Example-Based Adaptive Patch Sampling for Unsupervised Style Transfer
http://openaccess.thecvf.com/content_cvpr_2016/html/Frigo_Split_and_Match_CVPR_2016_paper.html
AUTHORS: Oriel Frigo, Neus Sabater, Julie Delon, Pierre Hellier
HIGHLIGHT: This paper presents a novel unsupervised method to transfer the style of an example image to a source image.
- 60, **TITLE:** Detection and Accurate Localization of Circular Fiducials Under Highly Challenging Conditions
http://openaccess.thecvf.com/content_cvpr_2016/html/Calvet_Detection_and_Accurate_CVPR_2016_paper.html
AUTHORS: Lilian Calvet, Pierre Gurdjos, Carsten Griwodz, Simone Gasparini
HIGHLIGHT: In this paper, we present a robust, highly accurate fiducial system, whose markers consist of concentric rings, along with its theoretical foundations.

- 61, TITLE: Scene Recognition With CNNs: Objects, Scales and Dataset Bias
http://openaccess.thecvf.com/content_cvpr_2016/html/Herranz_Scene_Recognition_With_CVPR_2016_paper.html
AUTHORS: Luis Herranz, Shuqiang Jiang, Xiangyang Li
HIGHLIGHT: In this paper we address two related problems: 1) scale induced dataset bias in multi-scale convolutional neural network (CNN) architectures, and 2) how to combine effectively scene-centric and object-centric knowledge (i.e. Places and ImageNet) in CNNs.
- 62, TITLE: Learning Action Maps of Large Environments via First-Person Vision
http://openaccess.thecvf.com/content_cvpr_2016/html/Rhinehart_Learning_Action_Maps_CVPR_2016_paper.html
AUTHORS: Nicholas Rhinehart, Kris M. Kitani
HIGHLIGHT: The method we describe enables functionality estimation in both large scenes where people have behaved, as well as novel scenes where no behaviors are available.
- 63, TITLE: Single-Image Crowd Counting via Multi-Column Convolutional Neural Network
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Single-Image_Crowd_Counting_CVPR_2016_paper.html
AUTHORS: Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, Yi Ma
HIGHLIGHT: This paper aims to develop a method that can accurately estimate the crowd count from an individual image with arbitrary crowd density and arbitrary perspective.
Since exiting crowd counting datasets do not adequately cover all the challenging situations considered in our work, we have collected and labelled a large new dataset that includes 1198 images with about 330,000 heads annotated.
- 64, TITLE: Shallow and Deep Convolutional Networks for Saliency Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Shallow_and_Deep_CVPR_2016_paper.html
AUTHORS: Junting Pan, Elisa Sayrol, Xavier Giro-i-Nieto, Kevin McGuinness, Noel E. O'Connor
HIGHLIGHT: Two designs are proposed: a shallow convnet trained from scratch, and a another deeper solution whose first three layers are adapted from another network trained for classification.
The recent publication of large datasets of saliency prediction has provided enough data to train end-to-end architectures that are both fast and accurate.
- 65, TITLE: Sample and Filter: Nonparametric Scene Parsing via Efficient Filtering
http://openaccess.thecvf.com/content_cvpr_2016/html/Najafi_Sample_and_Filter_CVPR_2016_paper.html
AUTHORS: Mohammad Najafi, Sarah Taghavi Namin, Mathieu Salzmann, Lars Petersson
HIGHLIGHT: In this paper, we introduce a nonparametric approach to scene parsing that follows a sample-and-filter strategy.
- 66, TITLE: DeLay: Robust Spatial Layout Estimation for Cluttered Indoor Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Dasgupta_DeLay_Robust_Spatial_CVPR_2016_paper.html
AUTHORS: Saumitro Dasgupta, Kuan Fang, Kevin Chen, Silvio Savarese
HIGHLIGHT: In this paper, we present a method that uses a fully convolutional neural network (FCNN) in conjunction with a novel optimization framework for generating layout estimates.
- 67, TITLE: A Text Detection System for Natural Scenes With Convolutional Feature Learning and Cascaded Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_A_Text_Detection_CVPR_2016_paper.html
AUTHORS: Siyu Zhu, Richard Zanibbi
HIGHLIGHT: We propose a system that finds text in natural scenes using a variety of cues.
- 68, TITLE: Reversible Recursive Instance-Level Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Liang_Reversible_Recursive_Instance-Level_CVPR_2016_paper.html
AUTHORS: Xiaodan Liang, Yunchao Wei, Xiaohui Shen, Zequn Jie, Jiashi Feng, Liang Lin, Shuicheng Yan
HIGHLIGHT: In this work, we propose a novel Reversible Recursive Instance-level Object Segmentation (R2-IOS) framework to address the challenging instance-level object segmentation task.
- 69, TITLE: Coherent Parametric Contours for Interactive Video Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Lu_Coherent_Parametric_Contours_CVPR_2016_paper.html
AUTHORS: Yao Lu, Xue Bai, Linda Shapiro, Jue Wang
HIGHLIGHT: We propose Coherent Parametric Contours, a novel video segmentation propagation framework that addresses all the above issues.
- 70, TITLE: Manifold SLIC: A Fast Method to Compute Content-Sensitive Superpixels
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Manifold_SLIC_A_CVPR_2016_paper.html
AUTHORS: Yong-Jin Liu, Cheng-Chi Yu, Min-Jing Yu, Ying He

HIGHLIGHT: In this paper, we extend SLIC to compute content-sensitive superpixels, i.e., small superpixels in content-dense regions (e.g., with high intensity or color variation) and large superpixels in content-sparse regions.

71, **TITLE:** Deep Saliency With Encoded Low Level Distance Map and High Level Features
http://openaccess.thecvf.com/content_cvpr_2016/html/Lee_Deep_Saliency_With_CVPR_2016_paper.html

AUTHORS: Gayoung Lee, Yu-Wing Tai, Junmo Kim

HIGHLIGHT: In this paper, we demonstrate that the hand-crafted features can provide complementary effects to enhance performance of saliency detection that utilizes only high level features.

72, **TITLE:** Instance-Level Segmentation for Autonomous Driving With Deep Densely Connected MRFs
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Instance-Level_Segmentation_for_CVPR_2016_paper.html

AUTHORS: Ziyu Zhang, Sanja Fidler, Raquel Urtasun

HIGHLIGHT: In this paper, we formulate the global labeling problem with a novel densely connected Markov random field and show how to encode various intuitive potentials in a way that is amenable to efficient mean field inference [Krahenbuhl et al., NIPS11].

73, **TITLE:** DHSNet: Deep Hierarchical Saliency Network for Salient Object Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_DHSNet_Deep_Hierarchical_CVPR_2016_paper.html

AUTHORS: Nian Liu, Junwei Han

HIGHLIGHT: In this work, we propose a novel end-to-end deep hierarchical saliency network (DHSNet) based on convolutional neural networks for detecting salient objects.

74, **TITLE:** Object Co-Segmentation via Graph Optimized-Flexible Manifold Ranking

http://openaccess.thecvf.com/content_cvpr_2016/html/Quan_Object_Co-Segmentation_via_CVPR_2016_paper.html

AUTHORS: Rong Quan, Junwei Han, Dingwen Zhang, Feiping Nie

HIGHLIGHT: To alleviate these limitations, we propose a novel two-stage co-segmentation framework, which introduces the weak background prior to establish a globally close-loop graph to represent the common object and union background separately.

75, **TITLE:** Primary Object Segmentation in Videos via Alternate Convex Optimization of Foreground and Background Distributions

http://openaccess.thecvf.com/content_cvpr_2016/html/Jang_Primary_Object_Segmentation_CVPR_2016_paper.html

AUTHORS: Won-Dong Jang, Chulwoo Lee, Chang-Su Kim

HIGHLIGHT: An unsupervised video object segmentation algorithm, which discovers a primary object in a video sequence automatically, is proposed in this work.

76, **TITLE:** Automatic Fence Segmentation in Videos of Dynamic Scenes

http://openaccess.thecvf.com/content_cvpr_2016/html/Yi_Automatic_Fence_Segmentation_CVPR_2016_paper.html

AUTHORS: Renjiao Yi, Jue Wang, Ping Tan

HIGHLIGHT: We present a fully automatic approach to detect and segment fence-like occluders from a video clip.

77, **TITLE:** Discovering the Physical Parts of an Articulated Object Class From Multiple Videos

http://openaccess.thecvf.com/content_cvpr_2016/html/Del_Pero_Discovering_the_Physical_CVPR_2016_paper.html

AUTHORS: Luca Del Pero, Susanna Ricco, Rahul Sukthankar, Vittorio Ferrari

HIGHLIGHT: We propose a motion-based method to discover the physical parts of an articulated object class (e.g. head/torso/leg of a horse) from multiple videos.

78, **TITLE:** A Benchmark Dataset and Evaluation Methodology for Video Object Segmentation

http://openaccess.thecvf.com/content_cvpr_2016/html/Perazzi_A_Benchmark_Dataset_CVPR_2016_paper.html

AUTHORS: Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, Alexander Sorkine-Hornung

HIGHLIGHT: In this work we present a new benchmark dataset and evaluation methodology for the area of video object segmentation.

79, **TITLE:** Learning Temporal Regularity in Video Sequences

http://openaccess.thecvf.com/content_cvpr_2016/html/Hasan_Learning_Temporal_Regularity_CVPR_2016_paper.html

AUTHORS: Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K. Roy-Chowdhury, Larry S. Davis

HIGHLIGHT: Specifically, we propose two methods that are built upon the autoencoders for their ability to work with little to no supervision.

80, **TITLE:** Bilateral Space Video Segmentation

http://openaccess.thecvf.com/content_cvpr_2016/html/Maerki_Bilateral_Space_Video_CVPR_2016_paper.html

AUTHORS: Nicolas Maerki, Federico Perazzi, Oliver Wang, Alexander Sorkine-Hornung

HIGHLIGHT: In this work, we propose a novel approach to video segmentation that operates in bilateral space.

81, TITLE: ReD-SFA: Relation Discovery Based Slow Feature Analysis for Trajectory Clustering

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_ReD-SFA_Relation_Discovery_CVPR_2016_paper.html

AUTHORS: Zhang Zhang, Kaiqi Huang, Tieniu Tan, Peipei Yang, Jun Li

HIGHLIGHT: In this paper, we proposed an approach, named Relation Discovery based Slow Feature Analysis (ReD-SFA), for feature learning and graph construction simultaneously.

82, TITLE: Training Region-Based Object Detectors With Online Hard Example Mining

http://openaccess.thecvf.com/content_cvpr_2016/html/Shrivastava_Training_Region-Based_Object_CVPR_2016_paper.html

AUTHORS: Abhinav Shrivastava, Abhinav Gupta, Ross Girshick

HIGHLIGHT: We present a simple yet surprisingly effective online hard example mining (OHEM) algorithm for training region-based ConvNet detectors.

83, TITLE: Deep Residual Learning for Image Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html

AUTHORS: Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun

HIGHLIGHT: We present a residual learning framework to ease the training of networks that are substantially deeper than those used previously.

84, TITLE: You Only Look Once: Unified, Real-Time Object Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Redmon_You_Only_Look_CVPR_2016_paper.html

AUTHORS: Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi

HIGHLIGHT: We present YOLO, a new approach to object detection.

85, TITLE: LocNet: Improving Localization Accuracy for Object Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Gidaris_LocNet_Improving_Localization_CVPR_2016_paper.html

AUTHORS: Spyros Gidaris, Nikos Komodakis

HIGHLIGHT: We propose a novel object localization methodology with the purpose of boosting the localization accuracy of state-of-the-art object detection systems.

86, TITLE: Sketch Me That Shoe

http://openaccess.thecvf.com/content_cvpr_2016/html/Yu_Sketch_Me_That_CVPR_2016_paper.html

AUTHORS: Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales, Chen-Change Loy

HIGHLIGHT: In this paper, for the first time, we address all these challenges, providing a step towards the capabilities that would underpin a commercial sketch-based image retrieval application.

We introduce a new database of 1,432 sketch-photo pairs from two categories with 32,000 fine-grained triplet ranking annotations.

87, TITLE: Deep Sliding Shapes for Amodal 3D Object Detection in RGB-D Images

http://openaccess.thecvf.com/content_cvpr_2016/html/Song_Deep_Sliding_Shapes_CVPR_2016_paper.html

AUTHORS: Shuran Song, Jianxiong Xiao

HIGHLIGHT: In our approach, we propose the first 3D Region Proposal Network (RPN) to learn objectness from geometric shapes and the first joint Object Recognition Network (ORN) to extract geometric features in 3D and color features in 2D.

88, TITLE: Object Detection From Video Tubelets With Convolutional Neural Networks

http://openaccess.thecvf.com/content_cvpr_2016/html/Kang_Object_Detection_From_CVPR_2016_paper.html

AUTHORS: Kai Kang, Wanli Ouyang, Hongsheng Li, Xiaogang Wang

HIGHLIGHT: In this work, we introduce a complete framework for the VID task based on still-image object detection and general object tracking.

89, TITLE: Learning With Side Information Through Modality Hallucination

http://openaccess.thecvf.com/content_cvpr_2016/html/Hoffman_Learning_With_Side_CVPR_2016_paper.html

AUTHORS: Judy Hoffman, Saurabh Gupta, Trevor Darrell

HIGHLIGHT: We present a modality hallucination architecture for training an RGB object detection model which incorporates depth side information at training time.

90, TITLE: Object-Proposal Evaluation Protocol is 'Gameable'

http://openaccess.thecvf.com/content_cvpr_2016/html/Chavali_Object-Proposal_Evaluation_Protocol_CVPR_2016_paper.html

AUTHORS: Neelima Chavali, Harsh Agrawal, Aroma Mahendru, Dhruv Batra
HIGHLIGHT: In this paper, we argue that the choice of using a partially annotated dataset for evaluation of object proposals is problematic -- as we demonstrate via a thought experiment, the evaluation protocol is 'gameable', in the sense that progress under this protocol does not necessarily correspond to a "better" category independent object proposal algorithm.

91, **TITLE:** HyperNet: Towards Accurate Region Proposal Generation and Joint Object Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Kong_HyperNet_Towards_Accurate_CVPR_2016_paper.html
AUTHORS: Tao Kong, Anbang Yao, Yurong Chen, Fuchun Sun
HIGHLIGHT: In this paper, we present a deep hierarchical network, namely HyperNet, for handling region proposal generation and object detection jointly.

92, **TITLE:** We Don't Need No Bounding-Boxes: Training Object Class Detectors Using Only Human Verification
http://openaccess.thecvf.com/content_cvpr_2016/html/Papadopoulos_We_Dont_Need_CVPR_2016_paper.html
AUTHORS: Dim P. Papadopoulos, Jasper R. R. Uijlings, Frank Keller, Vittorio Ferrari
HIGHLIGHT: We propose a new scheme for training object detectors which only requires annotators to verify bounding-boxes produced automatically by the learning algorithm.

93, **TITLE:** Factors in Finetuning Deep Model for Object Detection With Long-Tail Distribution
http://openaccess.thecvf.com/content_cvpr_2016/html/Ouyang_Factors_in_Finetuning_CVPR_2016_paper.html
AUTHORS: Wanli Ouyang, Xiaogang Wang, Cong Zhang, Xiaokang Yang
HIGHLIGHT: This paper investigates many factors that influence the performance in finetuning for object detection.

94, **TITLE:** Information-Driven Adaptive Structured-Light Scanners
http://openaccess.thecvf.com/content_cvpr_2016/html/Rosman_Information-Driven_Adaptive_Structured-Light_CVPR_2016_paper.html
AUTHORS: Guy Rosman, Daniela Rus, John W. Fisher III
HIGHLIGHT: In this paper we show how the same principles can be used as part of the 3D sensor.

95, **TITLE:** Simultaneous Optical Flow and Intensity Estimation From an Event Camera
http://openaccess.thecvf.com/content_cvpr_2016/html/Bardow_Simultaneous_Optical_Flow_CVPR_2016_paper.html
AUTHORS: Patrick Bardow, Andrew J. Davison, Stefan Leutenegger
HIGHLIGHT: Here, we propose, to the best of our knowledge, the first algorithm to simultaneously recover the motion field and brightness image, while the camera undergoes a generic motion through any scene.

96, **TITLE:** Macroscopic Interferometry: Rethinking Depth Estimation With Frequency-Domain Time-Of-Flight
http://openaccess.thecvf.com/content_cvpr_2016/html/Kadambi_Macroscopic_Interferometry_Rethinking_CVPR_2016_paper.html
AUTHORS: Achuta Kadambi, Jamie Schiel, Ramesh Raskar
HIGHLIGHT: This paper examines an alternative ToF architecture, inspired by micron-scale, microscopic interferometry, that relies only on frequency sampling: we refer to our proposed macroscopic technique as Frequency-Domain Time of Flight (FD-ToF).

97, **TITLE:** ASP Vision: Optically Computing the First Layer of Convolutional Neural Networks Using Angle Sensitive Pixels
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_ASP_Vision_Optically_CVPR_2016_paper.html
AUTHORS: Huaijin G. Chen, Suren Jayasuriya, Jiyue Yang, Judy Stephen, Sriram Sivaramakrishnan, Ashok Veeraraghavan, Alyosha Molnar
HIGHLIGHT: In this paper, we explore the energy savings of optically computing the first layer of CNNs.

98, **TITLE:** Computational Imaging for VLBI Image Reconstruction
http://openaccess.thecvf.com/content_cvpr_2016/html/Bouman_Computational_Imaging_for_CVPR_2016_paper.html
AUTHORS: Katherine L. Bouman, Michael D. Johnson, Daniel Zoran, Vincent L. Fish, Sheperd S. Doelman, William T. Freeman
HIGHLIGHT: In this paper we present a novel Bayesian approach for VLBI image reconstruction. We present this problem in a way that is accessible to members of the community, and provide a dataset website (vlbiimaging.csail.mit.edu) that facilitates controlled comparisons across algorithms.

99, **TITLE:** You Lead, We Exceed: Labor-Free Video Concept Learning by Jointly Exploiting Web Videos and Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Gan_You_Lead_We_CVPR_2016_paper.html
AUTHORS: Chuang Gan, Ting Yao, Kuiyuan Yang, Yi Yang, Tao Mei
HIGHLIGHT: In this paper, we propose a Lead--Exceed Neural Network (LENN), which reinforces the training on Web images and videos in a curriculum manner.

- 100, TITLE: Track and Segment: An Iterative Unsupervised Approach for Video Object Proposals
http://openaccess.thecvf.com/content_cvpr_2016/html/Xiao_Track_and_Segment_CVPR_2016_paper.html
AUTHORS: Fanyi Xiao, Yong Jae Lee
HIGHLIGHT: We present an unsupervised approach that generates a diverse, ranked set of bounding box and segmentation video object proposals---spatio-temporal tubes that localize the foreground objects---in an unannotated video.
- 101, TITLE: Beyond Local Search: Tracking Objects Everywhere With Instance-Specific Proposals
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Beyond_Local_Search_CVPR_2016_paper.html
AUTHORS: Gao Zhu, Fatih Porikli, Hongdong Li
HIGHLIGHT: Here, we present an object tracker that is not limited to a local search window and has ability to probe efficiently the entire frame.
- 102, TITLE: Groupwise Tracking of Crowded Similar-Appearance Targets From Low-Continuity Image Sequences
http://openaccess.thecvf.com/content_cvpr_2016/html/Yu_Groupwise_Tracking_of_CVPR_2016_paper.html
AUTHORS: Hongkai Yu, Youjie Zhou, Jeff Simmons, Craig P. Przybyla, Yuewei Lin, Xiaochuan Fan, Yang Mi, Song Wang
HIGHLIGHT: In this paper we propose a new groupwise method to explore the target group information and employ the within-group correlations for association and tracking.
- 103, TITLE: Social LSTM: Human Trajectory Prediction in Crowded Spaces
http://openaccess.thecvf.com/content_cvpr_2016/html/Alahi_Social_LSTM_Human_CVPR_2016_paper.html
AUTHORS: Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, Silvio Savarese
HIGHLIGHT: In our work, we propose a data-driven approach to learn these human-human interactions for predicting their future trajectories.
- 104, TITLE: What Players Do With the Ball: A Physically Constrained Interaction Modeling
http://openaccess.thecvf.com/content_cvpr_2016/html/Maksai_What_Players_Do_CVPR_2016_paper.html
AUTHORS: Andrii Maksai, Xinchao Wang, Pascal Fua
HIGHLIGHT: In this paper, we propose a generic and principled approach to modeling the interaction between the ball and the players while also imposing appropriate physical constraints on the ball's trajectory.
- 105, TITLE: Highlight Detection With Pairwise Deep Ranking for First-Person Video Summarization
http://openaccess.thecvf.com/content_cvpr_2016/html/Yao_Highlight_Detection_With_CVPR_2016_paper.html
AUTHORS: Ting Yao, Tao Mei, Yong Rui
HIGHLIGHT: Specifically, we propose a novel pairwise deep ranking model that employs deep learning techniques to learn the relationship between highlight and non-highlight video segments.
- 106, TITLE: Direct Prediction of 3D Body Poses From Motion Compensated Sequences
http://openaccess.thecvf.com/content_cvpr_2016/html/Tekin_Direct_Prediction_of_CVPR_2016_paper.html
AUTHORS: Bugra Tekin, Artem Rozantsev, Vincent Lepetit, Pascal Fua
HIGHLIGHT: We propose an efficient approach to exploiting motion information from consecutive frames of a video sequence to recover the 3D pose of people.
- 107, TITLE: Video2GIF: Automatic Generation of Animated GIFs From Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Gygli_Video2GIF_Automatic_Generation_CVPR_2016_paper.html
AUTHORS: Michael Gygli, Yale Song, Liangliang Cao
HIGHLIGHT: We introduce the novel problem of automatically generating animated GIFs from video.
- 108, TITLE: NTU RGB+D: A Large Scale Dataset for 3D Human Activity Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Shahroudy_NTU_RGBD_A_CVPR_2016_paper.html
AUTHORS: Amir Shahroudy, Jun Liu, Tian-Tsong Ng, Gang Wang
HIGHLIGHT: In this paper we introduce a large-scale dataset for RGB+D human action recognition with more than 56 thousand video samples and 4 million frames, collected from 40 distinct subjects.
- 109, TITLE: Progressively Parsing Interactional Objects for Fine Grained Action Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Ni_Progressively_Parsing_Interactional_CVPR_2016_paper.html
AUTHORS: Bingbing Ni, Xiaokang Yang, Shenghua Gao
HIGHLIGHT: In this work, we propose an end-to-end system based on recursive neural network to perform frame by frame interactional object parsing, which can alleviate the difficulty through an incremental manner.

- 110, TITLE: Hierarchical Recurrent Neural Encoder for Video Representation With Application to Captioning
http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Hierarchical_Recurrent_Neural_CVPR_2016_paper.html
AUTHORS: Pingbo Pan, Zhongwen Xu, Yi Yang, Fei Wu, Yueting Zhuang
HIGHLIGHT: In this paper, we propose a new approach, namely Hierarchical Recurrent Neural Encoder (HRNE), to exploit temporal information of videos.
- 111, TITLE: From Keyframes to Key Objects: Video Summarization by Representative Object Proposal Selection
http://openaccess.thecvf.com/content_cvpr_2016/html/Meng_From_Keyframes_to_CVPR_2016_paper.html
AUTHORS: Jingjing Meng, Hongxing Wang, Junsong Yuan, Yap-Peng Tan
HIGHLIGHT: We propose to summarize a video into a few key objects by selecting representative object proposals generated from video frames.
- 112, TITLE: Temporal Action Localization in Untrimmed Videos via Multi-Stage CNNs
http://openaccess.thecvf.com/content_cvpr_2016/html/Shou_Temporal_Action_Localization_CVPR_2016_paper.html
AUTHORS: Zheng Shou, Dongang Wang, Shih-Fu Chang
HIGHLIGHT: We propose a novel loss function for the localization network to explicitly consider temporal overlap and achieve high temporal localization accuracy.
- 113, TITLE: Summary Transfer: Exemplar-Based Subset Selection for Video Summarization
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Summary_Transfer_Exemplar-Based_CVPR_2016_paper.html
AUTHORS: Ke Zhang, Wei-Lun Chao, Fei Sha, Kristen Grauman
HIGHLIGHT: We propose a novel subset selection technique that leverages supervision in the form of human-created summaries to perform automatic keyframe-based video summarization.
- 114, TITLE: POD: Discovering Primary Objects in Videos Based on Evolutionary Refinement of Object Recurrence, Background, and Primary Object Models
http://openaccess.thecvf.com/content_cvpr_2016/html/Koh_POD_Discovering_Primary_CVPR_2016_paper.html
AUTHORS: Yeong Jun Koh, Won-Dong Jang, Chang-Su Kim
HIGHLIGHT: A primary object discovery (POD) algorithm for a video sequence is proposed in this work, which is capable of discovering a primary object, as well as identifying noisy frames that do not contain the object.
- 115, TITLE: What If We Do Not Have Multiple Videos of the Same Action? -- Video Action Localization Using Web Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Sultani_What_If_We_CVPR_2016_paper.html
AUTHORS: Waqas Sultani, Mubarak Shah
HIGHLIGHT: To obtain the most action representative proposal, we propose to reconstruct action proposals in the video by leveraging the action proposal in images.
- 116, TITLE: Beyond F-Formations: Determining Social Involvement in Free Standing Conversing Groups From Static Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Beyond_F-Formations_Determining_CVPR_2016_paper.html
AUTHORS: Lu Zhang, Hayley Hung
HIGHLIGHT: In this paper, we present the first attempt to analyse differing levels of social involvement in free standing conversing groups (or the so-called F-formations) from static images.
- 117, TITLE: DeepFashion: Powering Robust Clothes Recognition and Retrieval With Rich Annotations
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_DeepFashion_Powering_Robust_CVPR_2016_paper.html
AUTHORS: Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, Xiaoou Tang
HIGHLIGHT: To demonstrate the advantages of DeepFashion, we propose a new deep model, namely FashionNet, which learns clothing features by jointly predicting clothing attributes and landmarks.
- 118, TITLE: SketchNet: Sketch Classification With Web Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_SketchNet_Sketch_Classification_CVPR_2016_paper.html
AUTHORS: Hua Zhang, Si Liu, Changqing Zhang, Wenqi Ren, Rui Wang, Xiaochun Cao
HIGHLIGHT: In this study, we present a weakly supervised approach that discovers the discriminative structures of sketch images, given pairs of sketch images and web images.
- 119, TITLE: Embedding Label Structures for Fine-Grained Feature Representation
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Embedding_Label_Structures_CVPR_2016_paper.html
AUTHORS: Xiaofan Zhang, Feng Zhou, Yuanqing Lin, Shaoting Zhang

HIGHLIGHT: In this paper, we propose two main contributions to tackle this problem.

120, TITLE: Fine-Grained Image Classification by Exploring Bipartite-Graph Labels
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Fine-Grained_Image_Classification_CVPR_2016_paper.html
AUTHORS: Feng Zhou, Yuanqing Lin
HIGHLIGHT: This paper proposes a novel approach to exploit the rich relationships through bipartite-graph labels (BGL). To facilitate the study, we construct a new food benchmark dataset, which consists of 37,885 food images collected from 6 restaurants and totally 975 menus.

121, TITLE: Picking Deep Filter Responses for Fine-Grained Image Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Picking_Deep_Filter_CVPR_2016_paper.html
AUTHORS: Xiaopeng Zhang, Hongkai Xiong, Wengang Zhou, Weiyao Lin, Qi Tian
HIGHLIGHT: This paper proposes an automatic fine-grained recognition approach which is free of any object/part annotation at both training and testing stages.

122, TITLE: SPDA-CNN: Unifying Semantic Part Detection and Abstraction for Fine-Grained Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_SPDA-CNN_Unifying_Semantic_CVPR_2016_paper.html
AUTHORS: Han Zhang, Tao Xu, Mohamed Elhoseiny, Xiaolei Huang, Shaoting Zhang, Ahmed Elgammal, Dimitris Metaxas
HIGHLIGHT: In this paper, we propose a new CNN architecture that integrates semantic part detection and abstraction (SPDA-CNN) for fine-grained classification.

123, TITLE: Fine-Grained Categorization and Dataset Bootstrapping Using Deep Metric Learning With Humans in the Loop
http://openaccess.thecvf.com/content_cvpr_2016/html/Cui_Fine-Grained_Categorization_and_CVPR_2016_paper.html
AUTHORS: Yin Cui, Feng Zhou, Yuanqing Lin, Serge Belongie
HIGHLIGHT: In this work we propose a generic iterative framework for fine-grained categorization and dataset bootstrapping that handles these three challenges.

124, TITLE: Mining Discriminative Triplets of Patches for Fine-Grained Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Mining_Discriminative_Triplets_CVPR_2016_paper.html
AUTHORS: Yaming Wang, Jonghyun Choi, Vlad Morariu, Larry S. Davis
HIGHLIGHT: We describe a patch-based framework to address this problem.

125, TITLE: Part-Stacked CNN for Fine-Grained Visual Categorization
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Part-Stacked_CNN_for_CVPR_2016_paper.html
AUTHORS: Shaoli Huang, Zhe Xu, Dacheng Tao, Ya Zhang
HIGHLIGHT: In this paper, we propose a novel Part-Stacked CNN architecture that explicitly explains the fine-grained recognition process by modeling subtle differences from object parts.

126, TITLE: Learning Compact Binary Descriptors With Unsupervised Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Lin_Learning_Compact_Binary_CVPR_2016_paper.html
AUTHORS: Kevin Lin, Jiwen Lu, Chu-Song Chen, Jie Zhou
HIGHLIGHT: In this paper, we propose a new unsupervised deep learning approach called DeepBit to learn compact binary descriptor for efficient visual object matching.

127, TITLE: Solving Small-Piece Jigsaw Puzzles by Growing Consensus
http://openaccess.thecvf.com/content_cvpr_2016/html/Son_Solving_Small-Piece_Jigsaw_CVPR_2016_paper.html
AUTHORS: Kilho Son, daniel Moreno, James Hays, David B. Cooper
HIGHLIGHT: In this paper, we present a novel computational puzzle solver for square-piece image jigsaw puzzles with no prior information such as piece orientation, anchor pieces or resulting dimension of the puzzle.

128, TITLE: Pairwise Matching Through Max-Weight Bipartite Belief Propagation
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Pairwise_Matching_Through_CVPR_2016_paper.html
AUTHORS: Zhen Zhang, Qinfeng Shi, Julian McAuley, Wei Wei, Yanning Zhang, Anton van den Hengel
HIGHLIGHT: Here we revisit the use of graphical models for feature matching, and propose a belief propagation scheme which exhibits the following advantages: (1) we explicitly enforce one-to-one matching constraints; (2) we offer a tighter relaxation of the original cost function than previous graphical-model-based approaches; and (3) our sub-problems decompose into max-weight bipartite matching, which can be solved efficiently, leading to orders-of-magnitude reductions in execution time.

129, TITLE: Structured Feature Similarity With Explicit Feature Map

http://openaccess.thecvf.com/content_cvpr_2016/html/Kobayashi_Structured_Feature_Similarity_CVPR_2016_paper.html

AUTHORS: Takumi Kobayashi

HIGHLIGHT: In this paper, we propose a feature similarity measurement method based on the SSIM.

130, TITLE: Temporal Epipolar Regions

http://openaccess.thecvf.com/content_cvpr_2016/html/Dar_Temporal_Epipolar_Regions_CVPR_2016_paper.html

AUTHORS: Mor Dar, Yael Moses

HIGHLIGHT: In this paper we propose a novel method to predict locations of an approximately linear moving feature point, given a small subset of correspondences and the temporal order of image captures.

131, TITLE: Recurrent Attention Models for Depth-Based Person Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Haque_Recurrent_Attention_Models_CVPR_2016_paper.html

AUTHORS: Albert Haque, Alexandre Alahi, Li Fei-Fei

HIGHLIGHT: We present an attention-based model that reasons on human body shape and motion dynamics to identify individuals in the absence of RGB information, hence in the dark.

132, TITLE: Learning a Discriminative Null Space for Person Re-Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Learning_a_Discriminative_CVPR_2016_paper.html

AUTHORS: Li Zhang, Tao Xiang, Shaogang Gong

HIGHLIGHT: In this work, we propose to overcome the SSS problem in re-id distance metric learning by matching people in a discriminative null space of the training data.

133, TITLE: Learning Deep Feature Representations With Domain Guided Dropout for Person Re-Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Xiao_Learning_Deep_Feature_CVPR_2016_paper.html

AUTHORS: Tong Xiao, Hongsheng Li, Wanli Ouyang, Xiaogang Wang

HIGHLIGHT: In this work, we present a pipeline for learning deep feature representations from multiple domains with Convolutional Neural Networks (CNNs).

134, TITLE: How Far Are We From Solving Pedestrian Detection?

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_How_Far_Are_CVPR_2016_paper.html

AUTHORS: Shanshan Zhang, Rodrigo Benenson, Mohamed Omran, Jan Hosang, Bernt Schiele

HIGHLIGHT: We enable our analysis by creating a human baseline for pedestrian detection (over the Caltech dataset), and by manually clustering the recurrent errors of a top detector. Other than our in-depth analysis, we report top performance on the Caltech dataset, and provide a new sanitised set of training and test annotations.

135, TITLE: Similarity Learning With Spatial Constraints for Person Re-Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Similarity_Learning_With_CVPR_2016_paper.html

AUTHORS: Dapeng Chen, Zejian Yuan, Badong Chen, Nanning Zheng

HIGHLIGHT: We therefore learn a novel similarity function, which consists of multiple sub-similarity measurements with each taking in charge of a subregion.

136, TITLE: Sample-Specific SVM Learning for Person Re-Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Sample-Specific_SVM_Learning_CVPR_2016_paper.html

AUTHORS: Ying Zhang, Baohua Li, Huchuan Lu, Atshushi Irie, Xiang Ruan

HIGHLIGHT: In this paper, we formulate the person re-identification problem as an imbalanced classification problem and learn a classifier specifically for each pedestrian such that the matching model is highly tuned to the individual's appearance.

137, TITLE: Joint Learning of Single-Image and Cross-Image Representations for Person Re-Identification

http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Joint_Learning_of_CVPR_2016_paper.html

AUTHORS: Faqiang Wang, Wangmeng Zuo, Liang Lin, David Zhang, Lei Zhang

HIGHLIGHT: In this work, we exploit the connection between these two categories of methods, and propose a joint learning framework to unify SIR and CIR using convolutional neural network (CNN).

138, TITLE: A Multi-Level Contextual Model For Person Recognition in Photo Albums

http://openaccess.thecvf.com/content_cvpr_2016/html/Li_A_Multi-Level_Contextual_CVPR_2016_paper.html

AUTHORS: Haoxiang Li, Jonathan Brandt, Zhe Lin, Xiaohui Shen, Gang Hua

HIGHLIGHT: In this work, we present a new framework for person recognition in photo albums that exploits contextual cues at multiple levels, spanning individual persons, individual photos, and photo groups.

- 139, TITLE: Unsupervised Cross-Dataset Transfer Learning for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2016/html/Peng_Unsupervised_Cross-Dataset_Transfer_CVPR_2016_paper.html
AUTHORS: Peixi Peng, Tao Xiang, Yaowei Wang, Massimiliano Pontil, Shaogang Gong, Tiejun Huang, Yonghong Tian
HIGHLIGHT: Specifically, we present a multi-task dictionary learning method which is able to learn a dataset-shared but target-data-biased representation.
- 140, TITLE: Pedestrian Detection Inspired by Appearance Constancy and Shape Symmetry
http://openaccess.thecvf.com/content_cvpr_2016/html/Cao_Pedestrian_Detection_Inspired_CVPR_2016_paper.html
AUTHORS: Jiale Cao, Yanwei Pang, Xuelong Li
HIGHLIGHT: Inspired by some simple inherent attributes of pedestrians (i.e., appearance constancy and shape symmetry), we propose two new types of non-neighboring features (NNF): side-inner difference features (SIDF) and symmetrical similarity features (SSF).
- 141, TITLE: Recurrent Convolutional Network for Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2016/html/McLaughlin_Recurrent_Convolutional_Network_CVPR_2016_paper.html
AUTHORS: Niall McLaughlin, Jesus Martinez del Rincon, Paul Miller
HIGHLIGHT: In this paper we propose a novel recurrent neural network architecture for video-based person re-identification.
- 142, TITLE: Person Re-Identification by Multi-Channel Parts-Based CNN With Improved Triplet Loss Function
http://openaccess.thecvf.com/content_cvpr_2016/html/Cheng_Person_Re-Identification_by_CVPR_2016_paper.html
AUTHORS: De Cheng, Yihong Gong, Sanping Zhou, Jinjun Wang, Nanning Zheng
HIGHLIGHT: In this paper, we present a novel multi-channel parts-based convolutional neural network (CNN) model under the triplet framework for person re-identification.
- 143, TITLE: Top-Push Video-Based Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2016/html/You_Top-Push_Video-Based_Person_CVPR_2016_paper.html
AUTHORS: Jinjie You, Ancong Wu, Xiang Li, Wei-Shi Zheng
HIGHLIGHT: To solve this problem, we propose a top-push distance learning model (TDL), in which we integrate a top-push constraint, for matching video features of persons.
- 144, TITLE: Improving Person Re-Identification via Pose-Aware Multi-Shot Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Cho_Improving_Person_Re-Identification_CVPR_2016_paper.html
AUTHORS: Yeong-Jun Cho, Kuk-Jin Yoon
HIGHLIGHT: In this paper, we propose a novel framework for person re-identification by analyzing camera viewpoints and person poses, so-called Pose-aware Multi-shot Matching (PaMM), which robustly estimates target poses and efficiently conducts multi-shot matching based on the target pose information.
- 145, TITLE: Hierarchical Gaussian Descriptor for Person Re-Identification
http://openaccess.thecvf.com/content_cvpr_2016/html/Matsukawa_Hierarchical_Gaussian_Descriptor_CVPR_2016_paper.html
AUTHORS: Tetsu Matsukawa, Takahiro Okabe, Einoshin Suzuki, Yoichi Sato
HIGHLIGHT: In this paper, we present a novel descriptor based on a hierarchical distribution of pixel features.
- 146, TITLE: STCT: Sequentially Training Convolutional Networks for Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_STCT_Sequentially_Training_CVPR_2016_paper.html
AUTHORS: Lijun Wang, Wanli Ouyang, Xiaogang Wang, Huchuan Lu
HIGHLIGHT: In this paper, we propose a sequential training method for convolutional neural networks (CNNs) to effectively transfer pre-trained deep features for online applications.
- 147, TITLE: Determining Occlusions From Space and Time Image Reconstructions
http://openaccess.thecvf.com/content_cvpr_2016/html/Perez-Rua_Determining_Occlusions_From_CVPR_2016_paper.html
AUTHORS: Juan-Manuel Perez-Rua, Tomas Crivelli, Patrick Bouthemy, Patrick Perez
HIGHLIGHT: With this in mind, we propose a novel approach to occlusion detection where visibility or not of a point in next frame is formulated in terms of visual reconstruction.
- 148, TITLE: Online Multi-Object Tracking via Structural Constraint Event Aggregation
http://openaccess.thecvf.com/content_cvpr_2016/html/Yoon_Online_Multi-Object_Tracking_CVPR_2016_paper.html
AUTHORS: Ju Hong Yoon, Chang-Ryeol Lee, Ming-Hsuan Yang, Kuk-Jin Yoon
HIGHLIGHT: In this paper, we propose a new data association method that effectively exploits structural motion constraints in the presence of large camera motion.

- 149, TITLE: Staple: Complementary Learners for Real-Time Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Bertinetto_Staple_Complementary_Learners_CVPR_2016_paper.html
AUTHORS: Luca Bertinetto, Jack Valmadre, Stuart Golodetz, Ondrej Miksik, Philip H. S. Torr
HIGHLIGHT: In this paper, we show that a simple tracker combining complementary cues in a ridge regression framework can operate faster than 80 FPS and outperform not only all entries in the popular VOT14 competition, but also recent and far more sophisticated trackers according to multiple benchmarks.
- 150, TITLE: Robust Optical Flow Estimation of Double-Layer Images Under Transparency or Reflection
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Robust_Optical_Flow_CVPR_2016_paper.html
AUTHORS: Jiaolong Yang, Hongdong Li, Yuchao Dai, Robby T. Tan
HIGHLIGHT: In this paper, we propose a robust solution to this problem.
- 151, TITLE: Siamese Instance Search for Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Tao_Siamese_Instance_Search_CVPR_2016_paper.html
AUTHORS: Ran Tao, Efstratios Gavves, Arnold W.M. Smeulders
HIGHLIGHT: In this paper we present a tracker, which is radically different from state-of-the-art trackers: we apply no model updating, no occlusion detection, no combination of trackers, no geometric matching, and still deliver state-of-the-art tracking performance, as demonstrated on the popular online tracking benchmark (OTB) and six very challenging YouTube videos.
- 152, TITLE: Adaptive Decontamination of the Training Set: A Unified Formulation for Discriminative Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Danelljan_Adaptive_Decontamination_of_CVPR_2016_paper.html
AUTHORS: Martin Danelljan, Gustav Hager, Fahad Shahbaz Khan, Michael Felsberg
HIGHLIGHT: We propose a novel generic approach for alleviating the problem of corrupted training samples in tracking-by-detection frameworks.
- 153, TITLE: 3D Part-Based Sparse Tracker With Automatic Synchronization and Registration
http://openaccess.thecvf.com/content_cvpr_2016/html/Bibi_3D_Part-Based_Sparse_CVPR_2016_paper.html
AUTHORS: Adel Bibi, Tianzhu Zhang, Bernard Ghanem
HIGHLIGHT: In this paper, we present a part-based sparse tracker in a particle filter framework where both the motion and appearance model are formulated in 3D.
- 154, TITLE: Recurrently Target-Attending Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Cui_Recurrently_Target-Attending_Tracking_CVPR_2016_paper.html
AUTHORS: Zhen Cui, Shengtao Xiao, Jiashi Feng, Shuicheng Yan
HIGHLIGHT: To mitigate this problem, in this paper we propose a novel tracking method called Recurrently Target-attending Tracking (RTT).
- 155, TITLE: Structured Regression Gradient Boosting
http://openaccess.thecvf.com/content_cvpr_2016/html/Diego_Structured_Regression_Gradient_CVPR_2016_paper.html
AUTHORS: Ferran Diego, Fred A. Hamprecht
HIGHLIGHT: We propose a new way to train a structured output prediction model.
- 156, TITLE: Loss Functions for Top-k Error: Analysis and Insights
http://openaccess.thecvf.com/content_cvpr_2016/html/Lapin_Loss_Functions_for_CVPR_2016_paper.html
AUTHORS: Maksim Lapin, Matthias Hein, Bernt Schiele
HIGHLIGHT: In this paper, we provide an extensive comparison and evaluation of established multiclass methods comparing their top-k performance both from a practical as well as from a theoretical perspective.
- 157, TITLE: Metric Learning as Convex Combinations of Local Models With Generalization Guarantees
http://openaccess.thecvf.com/content_cvpr_2016/html/Zantedeschi_Metric_Learning_as_CVPR_2016_paper.html
AUTHORS: Valentina Zantedeschi, Remi Emonet, Marc Sebban
HIGHLIGHT: In this paper, we address these two issues by introducing a novel metric learning algorithm that linearly combines local models (C2LM).
- 158, TITLE: Efficient Training of Very Deep Neural Networks for Supervised Hashing
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Efficient_Training_of_CVPR_2016_paper.html
AUTHORS: Ziming Zhang, Yuting Chen, Venkatesh Saligrama
HIGHLIGHT: In this paper, we propose training very deep neural networks (DNNs) for supervised learning of hash codes.
- 159, TITLE: Information Bottleneck Learning Using Privileged Information for Visual Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Motiian_Information_Bottleneck_Learning_CVPR_2016_paper.html
AUTHORS: Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, Gianfranco Doretto
HIGHLIGHT: We explore the visual recognition problem from a main data view when an auxiliary data view is available during training.

160, TITLE: 3D Action Recognition From Novel Viewpoints
http://openaccess.thecvf.com/content_cvpr_2016/html/Rahmani_3D_Action_Recognition_CVPR_2016_paper.html
AUTHORS: Hossein Rahmani, Ajmal Mian
HIGHLIGHT: Therefore, we propose a method to generate this data by fitting synthetic 3D human models to real motion capture data and rendering the human poses from numerous viewpoints.

161, TITLE: 3D Shape Attributes
http://openaccess.thecvf.com/content_cvpr_2016/html/Fouhey_3D_Shape_Attributes_CVPR_2016_paper.html
AUTHORS: David F. Fouhey, Abhinav Gupta, Andrew Zisserman
HIGHLIGHT: In this paper we investigate 3D attributes as a means to understand the shape of an object in a single image. To this end, we make a number of contributions: (i) we introduce and define a set of 3D Shape attributes, including planarity, symmetry and occupied space; (ii) we show that such properties can be successfully inferred from a single image using a Convolutional Neural Network (CNN); (iii) we introduce a 143K image dataset of sculptures with 2197 works over 242 artists for training and evaluating the CNN; (iv) we show that the 3D attributes trained on this dataset generalize to images of other (non-sculpture) object classes; and furthermore (v) we show that the CNN also provides a shape embedding that can be used to match previously unseen sculptures largely independent of viewpoint.

162, TITLE: Three-Dimensional Object Detection and Layout Prediction Using Clouds of Oriented Gradients
http://openaccess.thecvf.com/content_cvpr_2016/html/Ren_Three-Dimensional_Object_Detection_CVPR_2016_paper.html
AUTHORS: Zhile Ren, Erik B. Sudderth
HIGHLIGHT: We propose a cloud of oriented gradient (COG) descriptor that links the 2D appearance and 3D pose of object categories, and thus accurately models how perspective projection affects perceived image boundaries.

163, TITLE: 3D Semantic Parsing of Large-Scale Indoor Spaces
http://openaccess.thecvf.com/content_cvpr_2016/html/Armeni_3D_Semantic_Parsing_CVPR_2016_paper.html
AUTHORS: Iro Armeni, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, Silvio Savarese
HIGHLIGHT: In this paper, we propose a method for semantic parsing the 3D point cloud of an entire building using a hierarchical approach: first, the raw data is parsed into semantically meaningful spaces (e.g. rooms, etc) that are aligned into a canonical reference coordinate system.

164, TITLE: Dense Human Body Correspondences Using Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Wei_Dense_Human_Body_CVPR_2016_paper.html
AUTHORS: Lingyu Wei, Qixing Huang, Duygu Ceylan, Etienne Vouga, Hao Li
HIGHLIGHT: We propose a deep learning approach for finding dense correspondences between 3D scans of people.

165, TITLE: Geometry-Informed Material Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/DeGol_Geometry-Informed_Material_Recognition_CVPR_2016_paper.html
AUTHORS: Joseph DeGol, Mani Golparvar-Fard, Derek Hoiem
HIGHLIGHT: Our goal is to recognize material categories using images and geometry information. We introduce a new dataset, GeoMat, which is the first to provide both image and geometry data in the form of: (i) training and testing patches that were extracted at different scales and perspectives from real world examples of each material category, and (ii) a large scale construction site scene that includes 160 images and over 800,000 hand labeled 3D points.

166, TITLE: Towards Open Set Deep Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Bendale_Towards_Open_Set_CVPR_2016_paper.html
AUTHORS: Abhijit Bendale, Terrance E. Boult
HIGHLIGHT: We present a methodology to adapt deep networks for open set recognition, by introducing a new model layer, OpenMax, which estimates the probability of an input being from an unknown class.

167, TITLE: What's Wrong With That Object? Identifying Images of Unusual Objects by Modelling the Detection Score Distribution
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Whats_Wrong_With_CVPR_2016_paper.html
AUTHORS: Peng Wang, Lingqiao Liu, Chunhua Shen, Zi Huang, Anton van den Hengel, Heng Tao Shen
HIGHLIGHT: To model these distributions we propose to use Gaussian Processes (GP) to construct two separate generative models, one for the "regular object" and the other for the "other objects".

- 168, TITLE: Large-Scale Location Recognition and the Geometric Burstiness Problem
http://openaccess.thecvf.com/content_cvpr_2016/html/Sattler_Large-Scale_Location_Recognition_CVPR_2016_paper.html
AUTHORS: Torsten Sattler, Michal Havlena, Konrad Schindler, Marc Pollefeys
HIGHLIGHT: We propose algorithms for detecting and handling geometric bursts.
- 169, TITLE: Regularity-Driven Facade Matching Between Aerial and Street Views
http://openaccess.thecvf.com/content_cvpr_2016/html/Wolff_Regularity-Driven_Facade_Matching_CVPR_2016_paper.html
AUTHORS: Mark Wolff, Robert T. Collins, Yanxi Liu
HIGHLIGHT: We present an approach for detecting and matching building facades between aerial view and street-view images.
- 170, TITLE: Do Computational Models Differ Systematically From Human Object Perception?
http://openaccess.thecvf.com/content_cvpr_2016/html/Pramod_Do_Computational_Models_CVPR_2016_paper.html
AUTHORS: R. T. Pramod, S. P. Arun
HIGHLIGHT: To answer this question we collected a large dataset of 26,675 perceived dissimilarity measurements from 2,801 visual objects across 269 human subjects, and used this dataset to train and test leading computational models.
- 171, TITLE: Contour Detection in Unstructured 3D Point Clouds
http://openaccess.thecvf.com/content_cvpr_2016/html/Hackel_Contour_Detection_in_CVPR_2016_paper.html
AUTHORS: Timo Hackel, Jan D. Wegner, Konrad Schindler
HIGHLIGHT: We describe a method to automatically detect contours, i.e. lines along which the surface orientation sharply changes, in large-scale outdoor point clouds.
- 172, TITLE: Unsupervised Learning of Edges
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Unsupervised_Learning_of_CVPR_2016_paper.html
AUTHORS: Yin Li, Manohar Paluri, James M. Rehg, Piotr Dollar
HIGHLIGHT: In this work we present a simple yet effective approach for training edge detectors without human supervision.
- 173, TITLE: Blind Image Deblurring Using Dark Channel Prior
http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Blind_Image_Deblurring_CVPR_2016_paper.html
AUTHORS: Jinsan Pan, Deqing Sun, Hanspeter Pfister, Ming-Hsuan Yang
HIGHLIGHT: We present a simple and effective blind image deblurring method based on the dark channel prior.
- 174, TITLE: Deeply-Recursive Convolutional Network for Image Super-Resolution
http://openaccess.thecvf.com/content_cvpr_2016/html/Kim_Deeply-Recursive_Convolutional_Network_CVPR_2016_paper.html
AUTHORS: Jiwon Kim, Jung Kwon Lee, Kyoung Mu Lee
HIGHLIGHT: We propose an image super-resolution method (SR) using a deeply-recursive convolutional network (DRCN).
- 175, TITLE: Accurate Image Super-Resolution Using Very Deep Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Kim_Accurate_Image_Super-Resolution_CVPR_2016_paper.html
AUTHORS: Jiwon Kim, Jung Kwon Lee, Kyoung Mu Lee
HIGHLIGHT: We present a highly accurate single image superresolution (SR) method.
- 176, TITLE: RAW Image Reconstruction Using a Self-Contained sRGB-JPEG Image With Only 64 KB Overhead
http://openaccess.thecvf.com/content_cvpr_2016/html/Nguyen_RAW_Image_Reconstruction_CVPR_2016_paper.html
AUTHORS: Rang M. H. Nguyen, Michael S. Brown
HIGHLIGHT: To address this issue, we present a method to encode the necessary metadata within an sRGB image to reconstruct a high-quality RAW image.
- 177, TITLE: Group MAD Competition - A New Methodology to Compare Objective Image Quality Models
http://openaccess.thecvf.com/content_cvpr_2016/html/Ma_Group_MAD_Competition_CVPR_2016_paper.html
AUTHORS: Kede Ma, Qingbo Wu, Zhou Wang, Zhengfang Duanmu, Hongwei Yong, Hongliang Li, Lei Zhang
HIGHLIGHT: Objective image quality assessment (IQA) models aim to automatically predict human visual perception of image quality and are of fundamental importance in the field of image processing and computer vision.
We first build a database that contains 4,744 source natural images, together with 94,880 distorted images created from them.
- 178, TITLE: Non-Local Image Dehazing
http://openaccess.thecvf.com/content_cvpr_2016/html/Berman_Non-Local_Image_DeHazing_CVPR_2016_paper.html
AUTHORS: Dana Berman, Tali treibitz, Shai Avidan
HIGHLIGHT: We, on the other hand, propose an algorithm based on a new, non-local prior.

- 179, TITLE: A Holistic Approach to Cross-Channel Image Noise Modeling and Its Application to Image Denoising
http://openaccess.thecvf.com/content_cvpr_2016/html/Nam_A_Holistic_Approach_CVPR_2016_paper.html
AUTHORS: Seonghyeon Nam, Youngbae Hwang, Yasuyuki Matsushita, Seon Joo Kim
HIGHLIGHT: We show the influence of the in-camera imaging pipeline on noise and propose a new noise model in the 3D RGB space to account for the color channel mix-ups.
- 180, TITLE: Multispectral Images Denoising by Intrinsic Tensor Sparsity Regularization
http://openaccess.thecvf.com/content_cvpr_2016/html/Xie_Multispectral_Images_Denoising_CVPR_2016_paper.html
AUTHORS: Qi Xie, Qian Zhao, Deyu Meng, Zongben Xu, Shuhang Gu, Wangmeng Zuo, Lei Zhang
HIGHLIGHT: In this paper, we propose a new tensor-based denoising approach by fully considering two intrinsic characteristics underlying an MSI, i.e., the global correlation along spectrum (GCS) and nonlocal self-similarity across space (NSS).
- 181, TITLE: A Comparative Study for Single Image Blind Deblurring
http://openaccess.thecvf.com/content_cvpr_2016/html/Lai_A_Comparative_Study_CVPR_2016_paper.html
AUTHORS: Wei-Sheng Lai, Jia-Bin Huang, Zhe Hu, Narendra Ahuja, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we aim to bridge this gap.
First, we collect a dataset of real blurred images and a dataset of synthetically blurred images.
- 182, TITLE: Spatiotemporal Bundle Adjustment for Dynamic 3D Reconstruction
http://openaccess.thecvf.com/content_cvpr_2016/html/Vo_Spatiotemporal_Bundle_Adjustment_CVPR_2016_paper.html
AUTHORS: Minh Vo, Srinivasa G. Narasimhan, Yaser Sheikh
HIGHLIGHT: In this paper, we present a spatiotemporal bundle adjustment approach that jointly optimizes four coupled sub-problems: estimating camera intrinsics and extrinsics, triangulating 3D static points, as well as subframe temporal alignment between cameras and estimating 3D trajectories of dynamic points.
- 183, TITLE: Inextensible Non-Rigid Shape-From-Motion by Second-Order Cone Programming
http://openaccess.thecvf.com/content_cvpr_2016/html/Chhatkuli_Inextensible_Non-Rigid_Shape-From-Motion_CVPR_2016_paper.html
AUTHORS: Ajad Chhatkuli, Daniel Pizarro, Toby Collins, Adrien Bartoli
HIGHLIGHT: We present a global and convex formulation for template-less 3D reconstruction of a deforming object with the perspective camera.
- 184, TITLE: Optimal Relative Pose With Unknown Correspondences
http://openaccess.thecvf.com/content_cvpr_2016/html/Fredriksson_Optimal_Relative_Pose_CVPR_2016_paper.html
AUTHORS: Johan Fredriksson, Viktor Larsson, Carl Olsson, Fredrik Kahl
HIGHLIGHT: In this paper, we go one step further and show that it is feasible to compute both the epipolar geometry and the correspondences at the same time based on geometry only.
- 185, TITLE: Homography Estimation From the Common Self-Polar Triangle of Separate Ellipses
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Homography_Estimation_From_CVPR_2016_paper.html
AUTHORS: Haifei Huang, Hui Zhang, Yiu-ming Cheung
HIGHLIGHT: In this paper, we address the problem of homography estimation from two separate ellipses.
- 186, TITLE: Heterogeneous Light Fields
http://openaccess.thecvf.com/content_cvpr_2016/html/Diebold_Heterogeneous_Light_Fields_CVPR_2016_paper.html
AUTHORS: Maximilian Diebold, Bernd Jahne, Alexander Gatto
HIGHLIGHT: Here, we introduce a modified structure tensor approach which improves depth estimation.
- 187, TITLE: A Consensus-Based Framework for Distributed Bundle Adjustment
http://openaccess.thecvf.com/content_cvpr_2016/html/Eriksson_A_Consensus-Based_Framework_CVPR_2016_paper.html
AUTHORS: Anders Eriksson, John Bastian, Tat-Jun Chin, Mats Isaksson
HIGHLIGHT: In this paper we study large-scale optimization problems in multi-view geometry, in particular the Bundle Adjustment problem.
- 188, TITLE: Globally Optimal Manhattan Frame Estimation in Real-Time
http://openaccess.thecvf.com/content_cvpr_2016/html/Joo_Globally_Optimal_Manhattan_CVPR_2016_paper.html
AUTHORS: Kyungdon Joo, Tae-Hyun Oh, Junsik Kim, In So Kweon
HIGHLIGHT: In this paper, we propose a novel bound computation method within an efficient measurement domain for MF estimation, i.e., the extended Gaussian image (EGI).

Given a set of surface normals, we pose a Manhattan Frame (MF) estimation problem as a consensus set maximization that maximizes the number of inliers over the rotation search space.

- 189, TITLE: Mirror Surface Reconstruction Under an Uncalibrated Camera
http://openaccess.thecvf.com/content_cvpr_2016/html/Han_Mirror_Surface_Reconstruction_CVPR_2016_paper.html
AUTHORS: Kai Han, Kwan-Yee K. Wong, Dirk Schnieders, Miaomiao Liu
HIGHLIGHT: This paper addresses the problem of mirror surface reconstruction, and a solution based on observing the reflections of a moving reference plane on the mirror surface is proposed.
- 190, TITLE: A Hole Filling Approach Based on Background Reconstruction for View Synthesis in 3D Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Luo_A_Hole_Filling_CVPR_2016_paper.html
AUTHORS: Guibo Luo, Yuesheng Zhu, Zhaotian Li, Liming Zhang
HIGHLIGHT: In this paper, a hole filling approach based on background reconstruction is proposed, in which the temporal correlation information in both the 2D video and its corresponding depth map are exploited to construct a background video.
- 191, TITLE: A Direct Least-Squares Solution to the PnP Problem With Unknown Focal Length
http://openaccess.thecvf.com/content_cvpr_2016/html/Zheng_A_Direct_Least-Squares_CVPR_2016_paper.html
AUTHORS: Yinqiang Zheng, Laurent Kneip
HIGHLIGHT: In this work, we propose a direct least-squares solution to the perspective-(n)-point (P(n)P) pose estimation problem of a partially calibrated camera, whose intrinsic parameters except the focal length are known.
- 192, TITLE: Efficient Intersection of Three Quadrics and Applications in Computer Vision
http://openaccess.thecvf.com/content_cvpr_2016/html/Kukelova_Efficient_Intersection_of_CVPR_2016_paper.html
AUTHORS: Zuzana Kukelova, Jan Heller, Andrew Fitzgibbon
HIGHLIGHT: In this paper, we present a new algorithm for finding all intersections of three quadrics.
- 193, TITLE: Using Spatial Order to Boost the Elimination of Incorrect Feature Matches
http://openaccess.thecvf.com/content_cvpr_2016/html/Talker_Using_Spatial_Order_CVPR_2016_paper.html
AUTHORS: Lior Talker, Yael Moses, Ilan Shimshoni
HIGHLIGHT: In this paper we propose an efficient method for estimating the number of correct matches without explicitly computing them.
- 194, TITLE: A Probabilistic Framework for Color-Based Point Set Registration
http://openaccess.thecvf.com/content_cvpr_2016/html/Danelljan_A_Probabilistic_Framework_CVPR_2016_paper.html
AUTHORS: Martin Danelljan, Giulia Meneghetti, Fahad Shahbaz Khan, Michael Felsberg
HIGHLIGHT: In this paper, we propose a probabilistic point set registration framework that exploits available color information associated with the points.
- 195, TITLE: Blind Image Deconvolution by Automatic Gradient Activation
http://openaccess.thecvf.com/content_cvpr_2016/html/Gong_Blind_Image_Deconvolution_CVPR_2016_paper.html
AUTHORS: Dong Gong, Mingkui Tan, Yanning Zhang, Anton van den Hengel, Qinfeng Shi
HIGHLIGHT: We thus introduce a gradient activation method to automatically select a subset of gradients of the latent image in a cutting-plane-based optimization scheme for kernel estimation.
- 196, TITLE: PSyCo: Manifold Span Reduction for Super Resolution
http://openaccess.thecvf.com/content_cvpr_2016/html/Perez-Pellitero_PSyCo_Manifold_Span_CVPR_2016_paper.html
AUTHORS: Eduardo Perez-Pellitero, Jordi Salvador, Javier Ruiz-Hidalgo, Bodo Rosenhahn
HIGHLIGHT: In this paper we present a novel regression-based SR algorithm that benefits from an extended knowledge of the structure of both manifolds.
We obtain the respective set of mirror-symmetry axes by means of a frequency analysis of the dihedral elements, and we use them to collapse the redundant variability through a modified symmetry distance.
- 197, TITLE: Parametric Object Motion From Blur
http://openaccess.thecvf.com/content_cvpr_2016/html/Gast_Parametric_Object_Motion_CVPR_2016_paper.html
AUTHORS: Jochen Gast, Anita Sellent, Stefan Roth
HIGHLIGHT: Drawing on the success of joint segmentation and parametric motion models in the context of optical flow estimation, we propose a parametric object motion model combined with a segmentation mask to exploit localized, non-uniform motion blur.
- 198, TITLE: Image Deblurring Using Smartphone Inertial Sensors

http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Image_Deblurring_Using_CVPR_2016_paper.html

AUTHORS: Zhe Hu, Lu Yuan, Stephen Lin, Ming-Hsuan Yang

HIGHLIGHT: In this paper, we identify the issues that plague smartphone inertial sensors and propose a solution that successfully utilizes the sensor readings for image deblurring.

199, TITLE: Seven Ways to Improve Example-Based Single Image Super Resolution

http://openaccess.thecvf.com/content_cvpr_2016/html/Timofte_Seven_Ways_to_CVPR_2016_paper.html

AUTHORS: Radu Timofte, Rasmus Rothe, Luc Van Gool

HIGHLIGHT: In this paper we present seven techniques that everybody should know to improve example-based single image super resolution (SR): 1) augmentation of data, 2) use of large dictionaries with efficient search structures, 3) cascading, 4) image self-similarities, 5) back projection refinement, 6) enhanced prediction by consistency check, and 7) context reasoning.

200, TITLE: Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network

http://openaccess.thecvf.com/content_cvpr_2016/html/Shi_Real-Time_Single_Image_CVPR_2016_paper.html

AUTHORS: Wenzhe Shi, Jose Caballero, Ferenc Huszar, Johannes Totz, Andrew P. Aitken, Rob Bishop, Daniel Rueckert, Zehan Wang

HIGHLIGHT: In this paper, we present the first convolutional neural network (CNN) capable of real-time SR of 1080p videos on a single K2 GPU.

201, TITLE: They Are Not Equally Reliable: Semantic Event Search Using Differentiated Concept Classifiers

http://openaccess.thecvf.com/content_cvpr_2016/html/Chang_They_Are_Not_CVPR_2016_paper.html

AUTHORS: Xiaojun Chang, Yao-Liang Yu, Yi Yang, Eric P. Xing

HIGHLIGHT: In this paper, we present a state-of-the-art event search system without any example videos.

202, TITLE: Going Deeper into First-Person Activity Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Ma_Going_Deeper_into_CVPR_2016_paper.html

AUTHORS: Minghuang Ma, Haoqi Fan, Kris M. Kitani

HIGHLIGHT: To integrate these ideas under one framework, we propose a twin stream network architecture, where one stream analyzes appearance information and the other stream analyzes motion information.

203, TITLE: Cascaded Interactional Targeting Network for Egocentric Video Analysis

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Cascaded_Interactional_Targeting_CVPR_2016_paper.html

AUTHORS: Yang Zhou, Bingbing Ni, Richang Hong, Xiaokang Yang, Qi Tian

HIGHLIGHT: This work aims to explicitly address these two issues via introducing a cascaded interactional targeting (i.e., infer both hand and active object regions) deep neural network.

204, TITLE: Fast Temporal Activity Proposals for Efficient Detection of Human Actions in Untrimmed Videos

http://openaccess.thecvf.com/content_cvpr_2016/html/Heilbron_Fast_Temporal_Activity_CVPR_2016_paper.html

AUTHORS: Fabian Caba Heilbron, Juan Carlos Niebles, Bernard Ghanem

HIGHLIGHT: In this paper, we introduce a proposal method that aims to recover temporal segments containing actions in untrimmed videos.

205, TITLE: Discriminative Hierarchical Rank Pooling for Activity Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Fernando_Discriminative_Hierarchical_Rank_CVPR_2016_paper.html

AUTHORS: Basura Fernando, Peter Anderson, Marcus Hutter, Stephen Gould

HIGHLIGHT: We present a method for jointly learning the video representation and activity classifier parameters.

206, TITLE: Convolutional Two-Stream Network Fusion for Video Action Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Feichtenhofer_Convolutional_Two-Stream_Network_CVPR_2016_paper.html

AUTHORS: Christoph Feichtenhofer, Axel Pinz, Andrew Zisserman

HIGHLIGHT: Recent applications of Convolutional Neural Networks (ConvNets) for human action recognition in videos have proposed different solutions for incorporating the appearance and motion information.

207, TITLE: Learning Activity Progression in LSTMs for Activity Detection and Early Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Ma_Learning_Activity_Progression_CVPR_2016_paper.html

AUTHORS: Shugao Ma, Leonid Sigal, Stan Sclaroff

HIGHLIGHT: In this work we improve training of temporal deep models to better learn activity progression for activity detection and early detection.

- 208, TITLE: VLAD3: Encoding Dynamics of Deep Features for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_VLAD3_Encoding_Dynamics_CVPR_2016_paper.html
AUTHORS: Yingwei Li, Weixin Li, Vijay Mahadevan, Nuno Vasconcelos
HIGHLIGHT: Here, we propose a representation, VLAD for Deep Dynamics (VLAD³), that accounts for different levels of video dynamics.
- 209, TITLE: A Multi-Stream Bi-Directional Recurrent Neural Network for Fine-Grained Action Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Singh_A_Multi-Stream_Bi-Directional_CVPR_2016_paper.html
AUTHORS: Bharat Singh, Tim K. Marks, Michael Jones, Oncel Tuzel, Ming Shao
HIGHLIGHT: We present a multi-stream bi-directional recurrent neural network for fine-grained action detection.
- 210, TITLE: A Hierarchical Deep Temporal Model for Group Activity Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Ibrahim_A_Hierarchical_Deep_CVPR_2016_paper.html
AUTHORS: Mostafa S. Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, Greg Mori
HIGHLIGHT: To make use of these observations, we present a 2-stage deep temporal model for the group activity recognition problem.
- 211, TITLE: A Hierarchical Pose-Based Approach to Complex Action Understanding Using Dictionaries of Actionlets and Motion Poselets
http://openaccess.thecvf.com/content_cvpr_2016/html/Lillo_A_Hierarchical_Pose-Based_CVPR_2016_paper.html
AUTHORS: Ivan Lillo, Juan Carlos Niebles, Alvaro Soto
HIGHLIGHT: In this paper, we introduce a new hierarchical model for human action recognition that is able to categorize complex actions performed in videos.
- 212, TITLE: A Key Volume Mining Deep Framework for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_A_Key_Volume_CVPR_2016_paper.html
AUTHORS: Wangjiang Zhu, Jie Hu, Gang Sun, Xudong Cao, Yu Qiao
HIGHLIGHT: To address this issue, we propose a key volume mining deep framework to identify key volumes and conduct classification simultaneously.
- 213, TITLE: Improved Hamming Distance Search Using Variable Length Substrings
http://openaccess.thecvf.com/content_cvpr_2016/html/Ong_Improved_Hamming_Distance_CVPR_2016_paper.html
AUTHORS: Eng-Jon Ong, Mirosław Bober
HIGHLIGHT: To this end, we propose a novel, unsupervised approach to thresholded search in Hamming space, supporting long codes (e.g. 512-bits) with a wide-range of Hamming distance radii.
- 214, TITLE: Shortlist Selection With Residual-Aware Distance Estimator for K-Nearest Neighbor Search
http://openaccess.thecvf.com/content_cvpr_2016/html/Heo_Shortlist_Selection_With_CVPR_2016_paper.html
AUTHORS: Jae-Pil Heo, Zhe Lin, Xiaohui Shen, Jonathan Brandt, Sung-eui Yoon
HIGHLIGHT: In this paper, we introduce a novel shortlist computation algorithm for approximate, high-dimensional nearest neighbor search.
- 215, TITLE: Supervised Quantization for Similarity Search
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Supervised_Quantization_for_CVPR_2016_paper.html
AUTHORS: Xiaojuan Wang, Ting Zhang, Guo-Jun Qi, Jinhui Tang, Jingdong Wang
HIGHLIGHT: In this paper, we address the problem of searching for semantically similar images from a large database.
- 216, TITLE: Efficient Large-Scale Approximate Nearest Neighbor Search on the GPU
http://openaccess.thecvf.com/content_cvpr_2016/html/Wieschollek_Efficient_Large-Scale_Approximate_CVPR_2016_paper.html
AUTHORS: Patrick Wieschollek, Oliver Wang, Alexander Sorkine-Hornung, Hendrik P. A. Lensch
HIGHLIGHT: We present a new approach for efficient approximate nearest neighbor (ANN) search in high dimensional spaces, extending the idea of Product Quantization.
- 217, TITLE: Collaborative Quantization for Cross-Modal Similarity Search
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Collaborative_Quantization_for_CVPR_2016_paper.html
AUTHORS: Ting Zhang, Jingdong Wang
HIGHLIGHT: This paper presents a compact coding solution for efficient search, with a focus on the quantization approach which has already shown the superior performance over the hashing solutions in the single-modal similarity search.
- 218, TITLE: Aggregating Image and Text Quantized Correlated Components

http://openaccess.thecvf.com/content_cvpr_2016/html/Tran_Aggregating_Image_and_CVPR_2016_paper.html
AUTHORS: Thi Quynh Nhi Tran, Herve Le Borgne, Michel Crucianu
HIGHLIGHT: Since they favor correlations against modality-specific information, these methods have shown some success in both cross-modal and bi-modal tasks.

219, TITLE: Efficient Indexing of Billion-Scale Datasets of Deep Descriptors
http://openaccess.thecvf.com/content_cvpr_2016/html/Babenko_Efficient_Indexing_of_CVPR_2016_paper.html
AUTHORS: Artem Babenko, Victor Lempitsky
HIGHLIGHT: In this paper, we introduce a new dataset of one billion descriptors based on DNNs and reveal the relative inefficiency of IMI-based indexing for such descriptors compared to SIFT data.

220, TITLE: Deep Supervised Hashing for Fast Image Retrieval
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Deep_Supervised_Hashing_CVPR_2016_paper.html
AUTHORS: Haomiao Liu, Ruiping Wang, Shiguang Shan, Xilin Chen
HIGHLIGHT: In this paper, we present a new hashing method to learn compact binary codes for highly efficient image retrieval on large-scale datasets.

221, TITLE: Efficient Large-Scale Similarity Search Using Matrix Factorization
http://openaccess.thecvf.com/content_cvpr_2016/html/Iscen_Efficient_Large-Scale_Similarity_CVPR_2016_paper.html
AUTHORS: Ahmet Iscen, Michael Rabbat, Teddy Furon
HIGHLIGHT: Our goal is to reduce the number of vector operations and memory for performing a search without sacrificing accuracy of the returned images.

222, TITLE: Incremental Object Discovery in Time-Varying Image Collections
http://openaccess.thecvf.com/content_cvpr_2016/html/Kontogianni_Incremental_Object_Discovery_CVPR_2016_paper.html
AUTHORS: Theodora Kontogianni, Markus Mathias, Bastian Leibe
HIGHLIGHT: Abstract In this paper, we address the problem of object discovery in time-varying, large-scale image collections.

223, TITLE: Detecting Migrating Birds at Night
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Detecting_Migrating_Birds_CVPR_2016_paper.html
AUTHORS: Jia-Bin Huang, Rich Caruana, Andrew Farnsworth, Steve Kelling, Narendra Ahuja
HIGHLIGHT: In this paper, we present a vision-based system for detecting migrating birds in flight at night.

224, TITLE: When Naive Bayes Nearest Neighbors Meet Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Kuzborskiy_When_Naive_Bayes_CVPR_2016_paper.html
AUTHORS: Ilja Kuzborskiy, Fabio Maria Carlucci, Barbara Caputo
HIGHLIGHT: This paper proposes a framework that addresses all these issues, thus bringing back NBNNs on the map.

225, TITLE: Traffic-Sign Detection and Classification in the Wild
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Traffic-Sign_Detection_and_CVPR_2016_paper.html
AUTHORS: Zhe Zhu, Dun Liang, Songhai Zhang, Xiaolei Huang, Baoli Li, Shimin Hu
HIGHLIGHT: The benchmark, source code and the CNN model introduced in this paper is publicly available. Firstly, we have created a large traffic-sign benchmark from 100000 Tencent Street View panoramas, going beyond previous benchmarks.

226, TITLE: Large Scale Semi-Supervised Object Detection Using Visual and Semantic Knowledge Transfer
http://openaccess.thecvf.com/content_cvpr_2016/html/Tang_Large_Scale_Semi-Supervised_CVPR_2016_paper.html
AUTHORS: Yuxing Tang, Josiah Wang, Boyang Gao, Emmanuel Dellandrea, Robert Gaizauskas, Liming Chen
HIGHLIGHT: We improve this previous work by incorporating knowledge about object similarities from visual and semantic domains during the transfer process.

227, TITLE: Exploit All the Layers: Fast and Accurate CNN Object Detector With Scale Dependent Pooling and Cascaded Rejection Classifiers
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Exploit_All_the_CVPR_2016_paper.html
AUTHORS: Fan Yang, Wongun Choi, Yuanqing Lin
HIGHLIGHT: In this paper, we investigate two new strategies to detect objects accurately and efficiently using deep convolutional neural network: 1) scale-dependent pooling and 2) layer-wise cascaded rejection classifiers.

228, TITLE: Dictionary Pair Classifier Driven Convolutional Neural Networks for Object Detection

- http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Dictionary_Pair_Classifier_CVPR_2016_paper.html
AUTHORS: Keze Wang, Liang Lin, Wangmeng Zuo, Shuhang Gu, Lei Zhang
HIGHLIGHT: In this paper, we propose a novel object detection system by unifying DPL with the convolutional feature learning.
- 229, TITLE: Monocular 3D Object Detection for Autonomous Driving
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Monocular_3D_Object_CVPR_2016_paper.html
AUTHORS: Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, Raquel Urtasun
HIGHLIGHT: The goal of this paper is to perform 3D object detection in single monocular images in the domain of autonomous driving.
- 230, TITLE: How Hard Can It Be? Estimating the Difficulty of Visual Search in an Image
http://openaccess.thecvf.com/content_cvpr_2016/html/Ionescu_How_Hard_Can_CVPR_2016_paper.html
AUTHORS: Radu Tudor Ionescu, Bogdan Alexe, Marius Leordeanu, Marius Popescu, Dim P. Papadopoulos, Vittorio Ferrari
HIGHLIGHT: We address the problem of estimating image difficulty defined as the human response time for solving a visual search task.
- 231, TITLE: Deep Relative Distance Learning: Tell the Difference Between Similar Vehicles
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Deep_Relative_Distance_CVPR_2016_paper.html
AUTHORS: Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, Tiejun Huang
HIGHLIGHT: We propose a Deep Relative Distance Learning (DRDL) method which exploits a two-branch deep convolutional network to project raw vehicle images into an Euclidean space where distance can be directly used to measure the similarity of arbitrary two vehicles.
- 232, TITLE: Eye Tracking for Everyone
http://openaccess.thecvf.com/content_cvpr_2016/html/Krafka_Eye_Tracking_for_CVPR_2016_paper.html
AUTHORS: Kyle Krafka, Aditya Khosla, Petr Kellnhofer, Harini Kannan, Suchendra Bhandarkar, Wojciech Matusik, Antonio Torralba
HIGHLIGHT: We tackle this problem by introducing GazeCapture, the first large-scale dataset for eye tracking, containing data from over 1450 people consisting of almost 2:5M frames.
- 233, TITLE: Efficient Globally Optimal 2D-To-3D Deformable Shape Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Lahner_Efficient_Globally_Optimal_CVPR_2016_paper.html
AUTHORS: Zorah Lahner, Emanuele Rodola, Frank R. Schmidt, Michael M. Bronstein, Daniel Cremers
HIGHLIGHT: We propose the first algorithm for non-rigid 2D-to-3D shape matching, where the input is a 2D query shape as well as a 3D target shape and the output is a continuous matching curve represented as a closed contour on the 3D shape.
- 234, TITLE: Ambiguity Helps: Classification With Disagreements in Crowdsourced Annotations
http://openaccess.thecvf.com/content_cvpr_2016/html/Sharmanska_Ambiguity_Helps_Classification_CVPR_2016_paper.html
AUTHORS: Viktoriia Sharmanska, Daniel Hernandez-Lobato, Jose Miguel Hernandez-Lobato, Novi Quadrianto
HIGHLIGHT: On the technical side, we propose a framework to incorporate annotation disagreements into the classifiers.
- 235, TITLE: A Task-Oriented Approach for Cost-Sensitive Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Mottaghi_A_Task-Oriented_Approach_CVPR_2016_paper.html
AUTHORS: Roozbeh Mottaghi, Hannaneh Hajishirzi, Ali Farhadi
HIGHLIGHT: In this paper, we propose a novel cost-sensitive task-oriented recognition method that is based on a combination of linguistic semantics and visual cues.
- 236, TITLE: Refining Architectures of Deep Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Shankar_Refining_Architectures_of_CVPR_2016_paper.html
AUTHORS: Sukrit Shankar, Duncan Robertson, Yani Ioannou, Antonio Criminisi, Roberto Cipolla
HIGHLIGHT: In this paper, we intend to answer this question and introduce a novel strategy that alters the architecture of a given CNN for a specified dataset, to potentially enhance the original accuracy while possibly reducing the model size.
- 237, TITLE: iLab-20M: A Large-Scale Controlled Object Dataset to Investigate Deep Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Borji_iLab-20M_A_Large-Scale_CVPR_2016_paper.html
AUTHORS: Ali Borji, Saeed Izadi, Laurent Itti
HIGHLIGHT: In this work, we introduce a large-scale synthetic dataset, which is freely and publicly available, and use it to answer several fundamental questions regarding selectivity and invariance properties of convolutional neural networks.

- 238, TITLE: Recursive Recurrent Nets With Attention Modeling for OCR in the Wild
http://openaccess.thecvf.com/content_cvpr_2016/html/Lee_Recursive_Recurrent_Nets_CVPR_2016_paper.html
AUTHORS: Chen-Yu Lee, Simon Osindero
HIGHLIGHT: We present recursive recurrent neural networks with attention modeling (R2AM) for lexicon-free optical character recognition in natural scene images.
- 239, TITLE: Deep Decision Network for Multi-Class Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Murthy_Deep_Decision_Network_CVPR_2016_paper.html
AUTHORS: Venkatesh N. Murthy, Vivek Singh, Terrence Chen, R. Manmatha, Dorin Comaniciu
HIGHLIGHT: In this paper, we present a novel Deep Decision Network (DDN) that provides an alternative approach towards building an efficient deep learning network.
- 240, TITLE: Less Is More: Zero-Shot Learning From Online Textual Documents With Noise Suppression
http://openaccess.thecvf.com/content_cvpr_2016/html/Qiao_Less_Is_More_CVPR_2016_paper.html
AUTHORS: Ruizhi Qiao, Lingqiao Liu, Chunhua Shen, Anton van den Hengel
HIGHLIGHT: More specifically, we propose an $l_{2,1}$ -norm based objective function which can simultaneously suppress the noisy signal in the text and learn a function to match the text document and visual features.
- 241, TITLE: Fast Algorithms for Linear and Kernel SVM+
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Fast_Algorithms_for_CVPR_2016_paper.html
AUTHORS: Wen Li, Dengxin Dai, Mingkui Tan, Dong Xu, Luc Van Gool
HIGHLIGHT: In this paper, we propose two efficient algorithms for solving the linear and kernel SVM+, respectively.
- 242, TITLE: Hierarchically Gated Deep Networks for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Qi_Hierarchically_Gated_Deep_CVPR_2016_paper.html
AUTHORS: Guo-Jun Qi
HIGHLIGHT: To address this challenge, we develop a novel paradigm of multi-scale deep network to model spatial contexts surrounding different pixels at various scales.
- 243, TITLE: Deep Structured Scene Parsing by Learning With Image Descriptions
http://openaccess.thecvf.com/content_cvpr_2016/html/Lin_Deep_Structured_Scene_CVPR_2016_paper.html
AUTHORS: Liang Lin, Guangrun Wang, Rui Zhang, Ruimao Zhang, Xiaodan Liang, Wangmeng Zuo
HIGHLIGHT: We propose a deep architecture consisting of two networks: i) a convolutional neural network (CNN) extracting the image representation for pixelwise object labeling and ii) a recursive neural network (RNN) discovering the hierarchical object structure and the inter-object relations.
- 244, TITLE: CNN-RNN: A Unified Framework for Multi-Label Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_CNN-RNN_A_Unified_CVPR_2016_paper.html
AUTHORS: Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, Wei Xu
HIGHLIGHT: In this paper, we utilize recurrent neural networks (RNNs) to address this problem.
- 245, TITLE: Walk and Learn: Facial Attribute Representation Learning From Egocentric Video and Contextual Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Walk_and_Learn_CVPR_2016_paper.html
AUTHORS: Jing Wang, Yu Cheng, Rogerio Schmidt Feris
HIGHLIGHT: We perform an extensive experimental analysis on wearable data and two standard benchmark datasets based on web images (LFWA and CelebA).
- 246, TITLE: CNN-N-Gram for Handwriting Word Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Poznanski_CNN-N-Gram_for_Handwriting_CVPR_2016_paper.html
AUTHORS: Arik Poznanski, Lior Wolf
HIGHLIGHT: CNN-N-Gram for Handwriting Word Recognition
- 247, TITLE: Synthetic Data for Text Localisation in Natural Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Gupta_Synthetic_Data_for_CVPR_2016_paper.html
AUTHORS: Ankush Gupta, Andrea Vedaldi, Andrew Zisserman
HIGHLIGHT: In this paper we introduce a new method for text detection in natural images.
- 248, TITLE: End-To-End People Detection in Crowded Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Stewart_End-To-End_People_Detection_CVPR_2016_paper.html

- AUTHORS: Russell Stewart, Mykhaylo Andriluka, Andrew Y. Ng
HIGHLIGHT: We propose a model that is based on decoding an image into a set of people detections.
- 249, TITLE: Real-Time Saliency Object Detection With a Minimum Spanning Tree
http://openaccess.thecvf.com/content_cvpr_2016/html/Tu_Real-Time_Saliency_Object_CVPR_2016_paper.html
AUTHORS: Wei-Chih Tu, Shengfeng He, Qingxiong Yang, Shao-Yi Chien
HIGHLIGHT: In this paper, we present a real-time saliency object detection system based on the minimum spanning tree.
- 250, TITLE: Local Background Enclosure for RGB-D Saliency Object Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Feng_Local_Background_Enclosure_CVPR_2016_paper.html
AUTHORS: David Feng, Nick Barnes, Shaodi You, Chris McCarthy
HIGHLIGHT: Here, we propose a novel RGB-D saliency feature.
- 251, TITLE: Adaptive Object Detection Using Adjacency and Zoom Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Lu_Adaptive_Object_Detection_CVPR_2016_paper.html
AUTHORS: Yongxi Lu, Tara Javidi, Svetlana Lazebnik
HIGHLIGHT: In this paper we propose to use a search strategy that adaptively directs computational resources to sub-regions likely to contain objects.
- 252, TITLE: Semantic Channels for Fast Pedestrian Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Costea_Semantic_Channels_for_CVPR_2016_paper.html
AUTHORS: Arthur Daniel Costea, Sergiu Nedevschi
HIGHLIGHT: In this work we propose a fast solution for achieving state of art results for both pedestrian detection and semantic segmentation.
- 253, TITLE: G-CNN: An Iterative Grid Based Object Detector
http://openaccess.thecvf.com/content_cvpr_2016/html/Najibi_G-CNN_An_Iterative_CVPR_2016_paper.html
AUTHORS: Mahyar Najibi, Mohammad Rastegari, Larry S. Davis
HIGHLIGHT: We introduce G-CNN, an object detection technique based on CNNs which works without proposal algorithms.
- 254, TITLE: Recurrent Face Aging
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Recurrent_Face_Aging_CVPR_2016_paper.html
AUTHORS: Wei Wang, Zhen Cui, Yan Yan, Jiashi Feng, Shuicheng Yan, Xiangbo Shu, Nicu Sebe
HIGHLIGHT: In this paper, we introduce a recurrent face aging (RFA) framework based on a recurrent neural network which can identify the ages of people from 0 to 80.
- 255, TITLE: Face2Face: Real-Time Face Capture and Reenactment of RGB Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Thies_Face2Face_Real-Time_Face_CVPR_2016_paper.html
AUTHORS: Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, Matthias Niessner
HIGHLIGHT: We present a novel approach for real-time facial reenactment of a monocular target video sequence (e.g., Youtube video).
- 256, TITLE: Self-Adaptive Matrix Completion for Heart Rate Estimation From Face Videos Under Realistic Conditions
http://openaccess.thecvf.com/content_cvpr_2016/html/Tulyakov_Self-Adaptive_Matrix_Completion_CVPR_2016_paper.html
AUTHORS: Sergey Tulyakov, Xavier Alameda-Pineda, Elisa Ricci, Lijun Yin, Jeffrey F. Cohn, Nicu Sebe
HIGHLIGHT: Opposite to previous approaches that estimate the HR by processing all the skin pixels inside a fixed region of interest, we introduce a strategy to dynamically select face regions useful for robust HR estimation.
- 257, TITLE: Visually Indicated Sounds
http://openaccess.thecvf.com/content_cvpr_2016/html/Owens_Visually_Indicated_Sounds_CVPR_2016_paper.html
AUTHORS: Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, William T. Freeman
HIGHLIGHT: In this paper, we propose the task of predicting what sound an object makes when struck as a way of studying physical interactions within a visual scene.
- 258, TITLE: Image Style Transfer Using Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Gatys_Image_Style_Transfer_CVPR_2016_paper.html
AUTHORS: Leon A. Gatys, Alexander S. Ecker, Matthias Bethge
HIGHLIGHT: We introduce A Neural Algorithm of Artistic Style that can separate and recombine the image content and style of natural images.

- 259, TITLE: Patch-Based Convolutional Neural Network for Whole Slide Tissue Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Hou_Patch-Based_Convolutional_Neural_CVPR_2016_paper.html
AUTHORS: Le Hou, Dimitris Samaras, Tahsin M. Kurc, Yi Gao, James E. Davis, Joel H. Saltz
HIGHLIGHT: We propose to train a decision fusion model to aggregate patch-level predictions given by patch-level CNNs, which to the best of our knowledge has not been shown before.
- 260, TITLE: Hedgehog Shape Priors for Multi-Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Isack_Hedgehog_Shape_Priors_CVPR_2016_paper.html
AUTHORS: Hossam Isack, Olga Veksler, Milan Sonka, Yuri Boykov
HIGHLIGHT: We propose a more general multi-object segmentation approach.
- 261, TITLE: Latent Variable Graphical Model Selection Using Harmonic Analysis: Applications to the Human Connectome Project (HCP)
http://openaccess.thecvf.com/content_cvpr_2016/html/Kim_Latent_Variable_Graphical_CVPR_2016_paper.html
AUTHORS: Won Hwa Kim, Hyunwoo J. Kim, Nagesh Adluru, Vikas Singh
HIGHLIGHT: A major goal of imaging studies such as the (ongoing) Human Connectome Project (HCP) is to characterize the structural network map of the human brain and identify its associations with covariates such as genotype, risk factors, and so on that correspond to an individual.
We provide a compelling set of scientific results on 500 scans from the recently released HCP data where our algorithm recovers highly interpretable and sparse conditional dependencies between brain connectivity pathways and well-known covariates.
- 262, TITLE: Simultaneous Estimation of Near IR BRDF and Fine-Scale Surface Geometry
http://openaccess.thecvf.com/content_cvpr_2016/html/Choe_Simultaneous_Estimation_of_CVPR_2016_paper.html
AUTHORS: Gyeongmin Choe, Srinivasa G. Narasimhan, In So Kweon
HIGHLIGHT: In this paper, we present an approach to simultaneously estimate NIR BRDF and fine-scale surface details by imaging materials under different IR lighting and viewing directions.
Our setup does not require complicated gantries or calibration and we present the first NIR dataset of 100 materials including a variety of fabrics (knits, weaves, cotton, satin, leather), and organic (skin, leaves, jute, trunk, fur) and inorganic materials (plastic, concrete, carpet).
- 263, TITLE: Do It Yourself Hyperspectral Imaging With Everyday Digital Cameras
http://openaccess.thecvf.com/content_cvpr_2016/html/Oh_Do_It_Yourself_CVPR_2016_paper.html
AUTHORS: Seoung Wug Oh, Michael S. Brown, Marc Pollefeys, Seon Joo Kim
HIGHLIGHT: In this paper, we present a framework for reconstructing hyperspectral images by using multiple consumer-level digital cameras.
- 264, TITLE: Automatic Content-Aware Color and Tone Stylization
http://openaccess.thecvf.com/content_cvpr_2016/html/Lee_Automatic_Content-Aware_Color_CVPR_2016_paper.html
AUTHORS: Joon-Young Lee, Kalyan Sunkavalli, Zhe Lin, Xiaohui Shen, In So Kweon
HIGHLIGHT: We introduce a new technique that automatically generates diverse, visually compelling stylizations for a photograph in an unsupervised manner.
- 265, TITLE: Combining Markov Random Fields and Convolutional Neural Networks for Image Synthesis
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Combining_Markov_Random_CVPR_2016_paper.html
AUTHORS: Chuan Li, Michael Wand
HIGHLIGHT: We apply the method to both photographic and non-photo-realistic (artwork) synthesis tasks.
- 266, TITLE: DCAN: Deep Contour-Aware Networks for Accurate Gland Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_DCAN_Deep_Contour-Aware_CVPR_2016_paper.html
AUTHORS: Hao Chen, Xiaojuan Qi, Lequan Yu, Pheng-Ann Heng
HIGHLIGHT: In this paper, we proposed an efficient deep contour-aware network (DCAN) to solve this challenging problem under a unified multi-task learning framework.
- 267, TITLE: Learning to Read Chest X-Rays: Recurrent Neural Cascade Model for Automated Image Annotation
http://openaccess.thecvf.com/content_cvpr_2016/html/Shin_Learning_to_Read_CVPR_2016_paper.html
AUTHORS: Hoo-Chang Shin, Kirk Roberts, Le Lu, Dina Demner-Fushman, Jianhua Yao, Ronald M. Summers
HIGHLIGHT: In this paper, we present a deep learning model to efficiently detect a disease from an image and annotate its contexts (e.g., location, severity and the affected organs).
- 268, TITLE: Conformal Surface Alignment With Optimal Mobius Search

- http://openaccess.thecvf.com/content_cvpr_2016/html/Le_Conformal_Surface_Alignment_CVPR_2016_paper.html
AUTHORS: Huu Le, Tat-Jun Chin, David Suter
HIGHLIGHT: We demonstrate the efficacy of our algorithm on data commonly used in computational conformal geometry.
- 269, TITLE: Coupled Harmonic Bases for Longitudinal Characterization of Brain Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Hwang_Coupled_Harmonic_Bases_CVPR_2016_paper.html
AUTHORS: Seong Jae Hwang, Nagesh Adluru, Maxwell D. Collins, Sathya N. Ravi, Barbara B. Bendlin, Sterling C. Johnson, Vikas Singh
HIGHLIGHT: We pose this important question in terms of the Laplacian of the connectivity graphs derived from various longitudinal or disease time points - quantifying its progression is then expressed in terms of coupling the harmonic bases of a full set of Laplacians.
- 270, TITLE: Automating Carotid Intima-Media Thickness Video Interpretation With Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Shin_Automating_Carotid_Intima-Media_CVPR_2016_paper.html
AUTHORS: Jae Shin, Nima Tajbakhsh, R. Todd Hurst, Christopher B. Kendall, Jianming Liang
HIGHLIGHT: To overcome this limitation, this paper presents a new system to automate CIMT video interpretation.
- 271, TITLE: Context Encoders: Feature Learning by Inpainting
http://openaccess.thecvf.com/content_cvpr_2016/html/Pathak_Context_Encoders_Feature_CVPR_2016_paper.html
AUTHORS: Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, Alexei A. Efros
HIGHLIGHT: We present an unsupervised visual feature learning algorithm driven by context-based pixel prediction.
- 272, TITLE: Comparative Deep Learning of Hybrid Representations for Image Recommendations
http://openaccess.thecvf.com/content_cvpr_2016/html/Lei_Comparative_Deep_Learning_CVPR_2016_paper.html
AUTHORS: Chenyi Lei, Dong Liu, Weiping Li, Zheng-Jun Zha, Houqiang Li
HIGHLIGHT: Such representations are termed hybrid and addressed via a deep learning approach in this paper.
- 273, TITLE: Fast ConvNets Using Group-Wise Brain Damage
http://openaccess.thecvf.com/content_cvpr_2016/html/Lebedev_Fast_ConvNets_Using_CVPR_2016_paper.html
AUTHORS: Vadim Lebedev, Victor Lempitsky
HIGHLIGHT: We investigate different ways to add group-wise pruning to the learning process, and show that several-fold speedups of convolutional layers can be attained using group-sparsity regularizers.
- 274, TITLE: Learning to Co-Generate Object Proposals With a Deep Structured Network
http://openaccess.thecvf.com/content_cvpr_2016/html/Hayder_Learning_to_Co-Generate_CVPR_2016_paper.html
AUTHORS: Zeeshan Hayder, Xuming He, Mathieu Salzmann
HIGHLIGHT: In this paper, we present an approach to co-generating object proposals in multiple images, thus leveraging the collective power of multiple object candidates.
- 275, TITLE: DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Moosavi-Dezfooli_DeepFool_A_Simple_CVPR_2016_paper.html
AUTHORS: Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Pascal Frossard
HIGHLIGHT: In this paper, we fill this gap and propose the DeepFool algorithm to efficiently compute perturbations that fool deep networks, and thus reliably quantify the robustness of these classifiers.
- 276, TITLE: Blockout: Dynamic Model Selection for Hierarchical Deep Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Murdock_Blockout_Dynamic_Model_CVPR_2016_paper.html
AUTHORS: Calvin Murdock, Zhen Li, Howard Zhou, Tom Duerig
HIGHLIGHT: Instead, we propose Blockout, a method for regularization and model selection that simultaneously learns both the model architecture and parameters.
- 277, TITLE: FireCaffe: Near-Linear Acceleration of Deep Neural Network Training on Compute Clusters
http://openaccess.thecvf.com/content_cvpr_2016/html/Iandola_FireCaffe_Near-Linear_Acceleration_CVPR_2016_paper.html
AUTHORS: Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Kurt Keutzer
HIGHLIGHT: In this paper we present FireCaffe, which successfully scales deep neural network training across a cluster of GPUs.
- 278, TITLE: MDL-CW: A Multimodal Deep Learning Framework With Cross Weights
http://openaccess.thecvf.com/content_cvpr_2016/html/Rastegar_MDL-CW_A_Multimodal_CVPR_2016_paper.html
AUTHORS: Sarah Rastegar, Mahdieh Soleymani, Hamid R. Rabiee, Seyed Mohsen Shojaee

HIGHLIGHT: In this paper, we propose a multimodal deep learning framework (MDL-CW) that exploits the cross weights between representation of modalities, and try to gradually learn interactions of the modalities in a deep network manner (from low to high level interactions).

279, **TITLE:** Structured Receptive Fields in CNNs
http://openaccess.thecvf.com/content_cvpr_2016/html/Jacobsen_Structured_Receptive_Fields_CVPR_2016_paper.html
AUTHORS: Jorn-Henrik Jacobsen, Jan van Gemert, Zhongyu Lou, Arnold W. M. Smeulders
HIGHLIGHT: We combine these ideas into structured receptive field networks, a model which has a fixed filter basis and yet retains the flexibility of CNNs.

280, **TITLE:** First Person Action Recognition Using Deep Learned Descriptors
http://openaccess.thecvf.com/content_cvpr_2016/html/Singh_First_Person_Action_CVPR_2016_paper.html
AUTHORS: Suriya Singh, Chetan Arora, C. V. Jawahar
HIGHLIGHT: We propose convolutional neural networks (CNNs) for end to end learning and classification of wearer's actions.

281, **TITLE:** Recognizing Micro-Actions and Reactions From Paired Egocentric Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Yonetani_Recognizing_Micro-Actions_and_CVPR_2016_paper.html
AUTHORS: Ryo Yonetani, Kris M. Kitani, Yoichi Sato
HIGHLIGHT: We aim to understand the dynamics of social interactions between two people by recognizing their actions and reactions using a head-mounted camera.
We also build a new dataset of dyadic (two-persons) interactions that comprises more than 1000 pairs of egocentric videos to enable systematic evaluations on the task of micro-action and reaction recognition.

282, **TITLE:** Mining 3D Key-Pose-Motifs for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Mining_3D_Key-Pose-Motifs_CVPR_2016_paper.html
AUTHORS: Chunyu Wang, Yizhou Wang, Alan L. Yuille
HIGHLIGHT: We propose an efficient algorithm to mine key-pose-motifs taking into account these probabilities.

283, **TITLE:** Predicting the Where and What of Actors and Actions Through Online Action Localization
http://openaccess.thecvf.com/content_cvpr_2016/html/Soomro_Predicting_the_Where_CVPR_2016_paper.html
AUTHORS: Khurram Soomro, Haroon Idrees, Mubarak Shah
HIGHLIGHT: This paper proposes a novel approach to tackle the challenging problem of 'online action localization' which entails predicting actions and their locations as they happen in a video.

284, **TITLE:** Actions ~ Transformations
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Actions_Transformations_CVPR_2016_paper.html
AUTHORS: Xiaolong Wang, Ali Farhadi, Abhinav Gupta
HIGHLIGHT: In this paper, we propose a novel representation for actions by modeling an action as a transformation which changes the state of the environment before the action happens (precondition) to the state after the action (effect).

285, **TITLE:** Visual Path Prediction in Complex Scenes With Crowded Moving Objects
http://openaccess.thecvf.com/content_cvpr_2016/html/Yoo_Visual_Path_Prediction_CVPR_2016_paper.html
AUTHORS: YoungJoon Yoo, Kimin Yun, Sangdoon Yun, JongHee Hong, Hawook Jeong, Jin Young Choi
HIGHLIGHT: This paper proposes a novel path prediction algorithm for progressing one step further than the existing works focusing on single target path prediction.

286, **TITLE:** End-To-End Learning of Action Detection From Frame Glimpses in Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Yeung_End-To-End_Learning_of_CVPR_2016_paper.html
AUTHORS: Serena Yeung, Olga Russakovsky, Greg Mori, Li Fei-Fei
HIGHLIGHT: In this work we introduce a fully end-to-end approach for action detection in videos that learns to directly predict the temporal bounds of actions.

287, **TITLE:** Action Recognition in Video Using Sparse Coding and Relative Features
http://openaccess.thecvf.com/content_cvpr_2016/html/Alfaro_Action_Recognition_in_CVPR_2016_paper.html
AUTHORS: Anali Alfaro, Domingo Mery, Alvaro Soto
HIGHLIGHT: This work presents an approach to category-based action recognition in video using sparse coding techniques.

288, **TITLE:** Improving Human Action Recognition by Non-Action Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Improving_Human_Action_CVPR_2016_paper.html

- AUTHORS: Yang Wang, Minh Hoai
HIGHLIGHT: In this paper we consider the task of recognizing human actions in realistic video where human actions are dominated by irrelevant factors.
- 289, TITLE: Actionness Estimation Using Hybrid Fully Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Actionness_Estimation_Using_CVPR_2016_paper.html
AUTHORS: Limin Wang, Yu Qiao, Xiaoou Tang, Luc Van Gool
HIGHLIGHT: This paper presents a new deep architecture for actionness estimation, called hybrid fully convolutional network (H-FCN), which is composed of appearance FCN (A-FCN) and motion FCN (M-FCN).
- 290, TITLE: Real-Time Action Recognition With Enhanced Motion Vector CNNs
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Real-Time_Action_Recognition_CVPR_2016_paper.html
AUTHORS: Bowen Zhang, Limin Wang, Zhe Wang, Yu Qiao, Hanli Wang
HIGHLIGHT: Our key insight for relieving this problem is that optical flow and motion vector are inherent correlated.
- 291, TITLE: Laplacian Patch-Based Image Synthesis
http://openaccess.thecvf.com/content_cvpr_2016/html/Lee_Laplacian_Patch-Based_Image_CVPR_2016_paper.html
AUTHORS: Joo Ho Lee, Inchang Choi, Min H. Kim
HIGHLIGHT: In this paper, we propose a patch-based synthesis using a Laplacian pyramid to improve searching correspondence with enhanced awareness of edge structures.
- 292, TITLE: Rain Streak Removal Using Layer Priors
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Rain_Streak_Removal_CVPR_2016_paper.html
AUTHORS: Yu Li, Robby T. Tan, Xiaojie Guo, Jiangbo Lu, Michael S. Brown
HIGHLIGHT: In this paper, we propose an effective method that uses simple patch-based priors for both the background and rain layers.
- 293, TITLE: Gradient-Domain Image Reconstruction Framework With Intensity-Range and Base-Structure Constraints
http://openaccess.thecvf.com/content_cvpr_2016/html/Shibata_Gradient-Domain_Image_Reconstruction_CVPR_2016_paper.html
AUTHORS: Takashi Shibata, Masayuki Tanaka, Masatoshi Okutomi
HIGHLIGHT: This paper presents a novel unified gradient-domain image reconstruction framework with intensity-range constraint and base-structure constraint.
- 294, TITLE: Removing Clouds and Recovering Ground Observations in Satellite Image Sequences via Temporally Contiguous Robust Matrix Completion
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Removing_Clouds_and_CVPR_2016_paper.html
AUTHORS: Jialei Wang, Peder A. Olsen, Andrew R. Conn, Aurelie C. Lozano
HIGHLIGHT: We consider the problem of removing and replacing clouds in satellite image sequences, which has a wide range of applications in remote sensing.
- 295, TITLE: D3: Deep Dual-Domain Based Fast Restoration of JPEG-Compressed Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_D3_Deep_Dual-Domain_CVPR_2016_paper.html
AUTHORS: Zhangyang Wang, Ding Liu, Shiyu Chang, Qing Ling, Yingzhen Yang, Thomas S. Huang
HIGHLIGHT: In this paper, we design a Deep Dual-Domain (D3) based fast restoration model to remove artifacts of JPEG compressed images.
- 296, TITLE: From Bows to Arrows: Rolling Shutter Rectification of Urban Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Rengarajan_From_Bows_to_CVPR_2016_paper.html
AUTHORS: Vijay Rengarajan, Ambasadram N. Rajagopalan, Rengarajan Aravind
HIGHLIGHT: We pose an optimization problem with line desirability costs based on straightness, angle, and length, to resolve the geometric ambiguities while estimating the camera motion based on a rotation-only model assuming known camera intrinsic matrix.
- 297, TITLE: A Weighted Variational Model for Simultaneous Reflectance and Illumination Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Fu_A_Weighted_Variational_CVPR_2016_paper.html
AUTHORS: Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, Xinghao Ding
HIGHLIGHT: We propose a weighted variational model to estimate both the reflectance and the illumination from an observed image.
- 298, TITLE: Visualizing and Understanding Deep Texture Representations

http://openaccess.thecvf.com/content_cvpr_2016/html/Lin_Visualizing_and_Understanding_CVPR_2016_paper.html

AUTHORS: Tsung-Yu Lin, Subhansu Maji

HIGHLIGHT: Based on recent work [13, 28] we propose a technique to visualize pre-images, providing a means for understanding categorical properties that are captured by these representations.

299, TITLE: Robust Kernel Estimation With Outliers Handling for Image Deblurring

http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Robust_Kernel_Estimation_CVPR_2016_paper.html

AUTHORS: Jinshan Pan, Zhouchen Lin, Zhixun Su, Ming-Hsuan Yang

HIGHLIGHT: In this paper, we present an algorithm to address this problem by exploiting reliable edges and removing outliers in the intermediate latent images, thereby estimating blur kernels robustly.

300, TITLE: Online Collaborative Learning for Open-Vocabulary Visual Classifiers

http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Online_Collaborative_Learning_CVPR_2016_paper.html

AUTHORS: Hanwang Zhang, Xindi Shang, Wenzhuo Yang, Huan Xu, Huanbo Luan, Tat-Seng Chua

HIGHLIGHT: In this paper, we propose a novel online learning paradigm towards this challenging task.

301, TITLE: Rethinking the Inception Architecture for Computer Vision

http://openaccess.thecvf.com/content_cvpr_2016/html/Szegedy_Rethinking_the_Inception_CVPR_2016_paper.html

AUTHORS: Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, Zbigniew Wojna

HIGHLIGHT: Here we are exploring ways to scale up networks in ways that aim at utilizing the added computation as efficiently as possible.

302, TITLE: Cross Modal Distillation for Supervision Transfer

http://openaccess.thecvf.com/content_cvpr_2016/html/Gupta_Cross_Modal_Distillation_CVPR_2016_paper.html

AUTHORS: Saurabh Gupta, Judy Hoffman, Jitendra Malik

HIGHLIGHT: In this work we propose a technique that transfers supervision between images from different modalities.

303, TITLE: Efficient Point Process Inference for Large-Scale Object Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Pham_Efficient_Point_Process_CVPR_2016_paper.html

AUTHORS: Trung T. Pham, Seyed Hamid Rezatofighi, Ian Reid, Tat-Jun Chin

HIGHLIGHT: We propose an efficient point process inference for large-scale object detection using discrete energy minimization.

304, TITLE: Weakly Supervised Deep Detection Networks

http://openaccess.thecvf.com/content_cvpr_2016/html/Bilen_Weakly_Supervised_Deep_CVPR_2016_paper.html

AUTHORS: Hakan Bilen, Andrea Vedaldi

HIGHLIGHT: In this paper, we address this problem by exploiting the power of deep convolutional neural networks pre-trained on large-scale image-level classification tasks.

305, TITLE: BORDER: An Oriented Rectangles Approach to Texture-Less Object Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Chan_BORDER_An_Oriented_CVPR_2016_paper.html

AUTHORS: Jacob Chan, Jimmy Addison Lee, Qian Kemao

HIGHLIGHT: We correspondingly introduce a modified line-segment detection technique termed Linelets to stabilize keypoint repeatability in homogenous conditions.

306, TITLE: Active Image Segmentation Propagation

http://openaccess.thecvf.com/content_cvpr_2016/html/Jain_Active_Image_Segmentation_CVPR_2016_paper.html

AUTHORS: Suyog Dutt Jain, Kristen Grauman

HIGHLIGHT: We propose a semi-automatic method to obtain foreground object masks for a large set of related images.

307, TITLE: Inside-Outside Net: Detecting Objects in Context With Skip Pooling and Recurrent Neural Networks

http://openaccess.thecvf.com/content_cvpr_2016/html/Bell_Inside-Outside_Net_Detecting_CVPR_2016_paper.html

AUTHORS: Sean Bell, C. Lawrence Zitnick, Kavita Bala, Ross Girshick

HIGHLIGHT: In this paper we present the Inside-Outside Net (ION), an object detector that exploits information both inside and outside the region of interest.

308, TITLE: RIFD-CNN: Rotation-Invariant and Fisher Discriminative Convolutional Neural Networks for Object Detection

http://openaccess.thecvf.com/content_cvpr_2016/html/Cheng_RIFD-CNN_Rotation-Invariant_and_CVPR_2016_paper.html

AUTHORS: Gong Cheng, Peicheng Zhou, Junwei Han

HIGHLIGHT: To address these problems, this paper proposes a novel and effective method to learn a rotation-invariant and Fisher discriminative CNN (RIFD-CNN) model.

309, **TITLE:** Reinforcement Learning for Visual Object Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Mathe_Reinforcement_Learning_for_CVPR_2016_paper.html
AUTHORS: Stefan Mathe, Aleksis Pirinen, Cristian Sminchisescu
HIGHLIGHT: In this paper we present formally rigorous sequential models that accumulate evidence collected at a small set of image locations in order to detect visual objects effectively.

310, **TITLE:** Detecting Repeating Objects Using Patch Correlation Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Huberman_Detecting_Repeating_Objects_CVPR_2016_paper.html
AUTHORS: Inbar Huberman, Raanan Fattal
HIGHLIGHT: In this paper we describe a new method for detecting and counting a repeating object in an image.

311, **TITLE:** Analyzing Classifiers: Fisher Vectors and Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Bach_Analyzing_Classifiers_Fisher_CVPR_2016_paper.html
AUTHORS: Sebastian Bach, Alexander Binder, Gregoire Montavon, Klaus-Robert Muller, Wojciech Samek
HIGHLIGHT: Recently, a principled technique, Layer-wise Relevance Propagation (LRP), has been developed in order to better comprehend the inherent structured reasoning of complex nonlinear classification models such as Bag of Feature models or DNNs.

312, **TITLE:** Learning Deep Features for Discriminative Localization
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Learning_Deep_Features_CVPR_2016_paper.html
AUTHORS: Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, Antonio Torralba
HIGHLIGHT: In this work, we revisit the global average pooling layer proposed in [13], and shed light on how it explicitly enables the convolutional neural network (CNN) to have remarkable localization ability despite being trained on image-level labels.

313, **TITLE:** Seeing Through the Human Reporting Bias: Visual Classifiers From Noisy Human-Centric Labels
http://openaccess.thecvf.com/content_cvpr_2016/html/Misra_Seeing_Through_the_CVPR_2016_paper.html
AUTHORS: Ishan Misra, C. Lawrence Zitnick, Margaret Mitchell, Ross Girshick
HIGHLIGHT: We propose an algorithm to decouple the human reporting bias from the correct visually grounded labels.

314, **TITLE:** Learning Aligned Cross-Modal Representations From Weakly Aligned Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Castrejon_Learning_Aligned_Cross-Modal_CVPR_2016_paper.html
AUTHORS: Lluís Castrejon, Yusuf Aytar, Carl Vondrick, Hamed Pirsiavash, Antonio Torralba
HIGHLIGHT: In this paper, we investigate how to learn cross-modal scene representations that transfer across modalities. To study this problem, we introduce a new cross-modal scene dataset.

315, **TITLE:** A Probabilistic Collaborative Representation Based Approach for Pattern Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Cai_A_Probabilistic_Collaborative_CVPR_2016_paper.html
AUTHORS: Sijia Cai, Lei Zhang, Wangmeng Zuo, Xiangchu Feng
HIGHLIGHT: In this paper we propose a probabilistic collaborative representation framework, where the probability that a test sample belongs to the collaborative subspace of all classes can be well defined and computed.

316, **TITLE:** Learning Structured Inference Neural Networks With Label Relations
http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Learning_Structured_Inference_CVPR_2016_paper.html
AUTHORS: Hexiang Hu, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, Greg Mori
HIGHLIGHT: In this paper, we exploit this rich information with a state-of-art deep learning framework, and propose a generic structured model that leverages diverse label relations to improve image classification performance.

317, **TITLE:** Discriminative Multi-Modal Feature Fusion for RGBD Indoor Scene Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Discriminative_Multi-Modal_Feature_CVPR_2016_paper.html
AUTHORS: Hongyuan Zhu, Jean-Baptiste Weibel, Shijian Lu
HIGHLIGHT: Inspired by some recent work on RGBD object recognition using multi-modal feature fusion, we introduce a novel discriminative multi-modal fusion framework for rgb-d scene recognition for the first time which simultaneously considers the inter- and intra-modality correlation for all samples and meanwhile regularizing the learned features to be discriminative and compact.

318, **TITLE:** Conditional Graphical Lasso for Multi-Label Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Conditional_Graphical_Lasso_CVPR_2016_paper.html
AUTHORS: Qiang Li, Maoying Qiao, Wei Bian, Dacheng Tao

- HIGHLIGHT:** In this paper, we develop conditional graphical Lasso (CGL) to handle these challenges.
- 319, **TITLE:** Region Ranking SVM for Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Wei_Region_Ranking_SVM_CVPR_2016_paper.html
AUTHORS: Zijun Wei, Minh Hoai
HIGHLIGHT: In this paper we propose Region Ranking SVM(RRSVM), a novel method for pooling local information from multiple regions.
- 320, **TITLE:** Predicting Motivations of Actions by Leveraging Text
http://openaccess.thecvf.com/content_cvpr_2016/html/Vondrick_Predicting_Motivations_of_CVPR_2016_paper.html
AUTHORS: Carl Vondrick, Deniz Oktay, Hamed Pirsiavash, Antonio Torralba
HIGHLIGHT: In this paper, we introduce the problem of predicting why a person has performed an action in images. To study this problem, we introduce a new dataset of people performing actions annotated with likely motivations.
- 321, **TITLE:** BoxCars: 3D Boxes as CNN Input for Improved Fine-Grained Vehicle Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Sochor_BoxCars_3D_Boxes_CVPR_2016_paper.html
AUTHORS: Jakub Sochor, Adam Herout, Jiri Havel
HIGHLIGHT: Our contribution is showing that extracting additional data from the video stream - besides the vehicle image itself - and feeding it into the deep convolutional neural network boosts the recognition performance considerably.
- 322, **TITLE:** Highway Vehicle Counting in Compressed Domain
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Highway_Vehicle_Counting_CVPR_2016_paper.html
AUTHORS: Xu Liu, Zilei Wang, Jiashi Feng, Hongsheng Xi
HIGHLIGHT: This paper presents a highway vehicle counting method in compressed domain, aiming at achieving acceptable estimation performance approaching the pixel-domain methods.
- 323, **TITLE:** Camera Calibration From Periodic Motion of a Pedestrian
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Camera_Calibration_From_CVPR_2016_paper.html
AUTHORS: Shiyao Huang, Xianghua Ying, Jiangpeng Rong, Zeyu Shang, Hongbin Zha
HIGHLIGHT: In this paper, we propose a novel camera calibration method based on recovering the three orthogonal vanishing points (TOVPs), just using an image sequence of a pedestrian walking in a straight line, without any assumption of scenes or motions, e.g., control points with known 3D coordinates, parallel or perpendicular lines, non-natural or pre-designed special human motions, as often necessary in previous methods.
- 324, **TITLE:** Dynamic Image Networks for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Bilen_Dynamic_Image_Networks_CVPR_2016_paper.html
AUTHORS: Hakan Bilen, Basura Fernando, Efstratios Gavves, Andrea Vedaldi, Stephen Gould
HIGHLIGHT: We introduce the concept of dynamic image, a novel compact representation of videos useful for video analysis especially when convolutional neural networks (CNNs) are used.
- 325, **TITLE:** Detecting Events and Key Actors in Multi-Person Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Ramanathan_Detecting_Events_and_CVPR_2016_paper.html
AUTHORS: Vignesh Ramanathan, Jonathan Huang, Sami Abu-El-Haija, Alexander Gorban, Kevin Murphy, Li Fei-Fei
HIGHLIGHT: In this paper, we propose a model which learns to detect events in such videos while automatically "attending" to the people responsible for the event.
Since most video datasets with multiple people are restricted to a small number of videos, we also collected a new basketball dataset comprising 257 basketball games with 14K event annotations corresponding to 11 event classes.
- 326, **TITLE:** Regularizing Long Short Term Memory With 3D Human-Skeleton Sequences for Action Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Mahasseni_Regularizing_Long_Short_CVPR_2016_paper.html
AUTHORS: Behrooz Mahasseni, Sinisa Todorovic
HIGHLIGHT: This paper argues that large-scale action recognition in video can be greatly improved by providing an additional modality in training data -- namely, 3D human-skeleton sequences -- aimed at complementing poorly represented or missing features of human actions in the training videos.
- 327, **TITLE:** Personalizing Human Video Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Charles_Personalizing_Human_Video_CVPR_2016_paper.html
AUTHORS: James Charles, Tomas Pfister, Derek Mgee, David Hogg, Andrew Zisserman
HIGHLIGHT: We propose a personalized ConvNet pose estimator that automatically adapts itself to the uniqueness of a person's appearance to improve pose estimation in long videos.

- 328, TITLE: End-To-End Learning of Deformable Mixture of Parts and Deep Convolutional Neural Networks for Human Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_End-To-End_Learning_of_CVPR_2016_paper.html
AUTHORS: Wei Yang, Wanli Ouyang, Hongsheng Li, Xiaogang Wang
HIGHLIGHT: In this paper, we propose a novel end-to-end framework for human pose estimation that combines DCNNs with the expressive deformable mixture of parts.
- 329, TITLE: Actor-Action Semantic Segmentation With Grouping Process Models
http://openaccess.thecvf.com/content_cvpr_2016/html/Xu_Actor-Action_Semantic_Segmentation_CVPR_2016_paper.html
AUTHORS: Chenliang Xu, Jason J. Corso
HIGHLIGHT: We propose a new model that combines the labeling CRF with a supervoxel hierarchy, where supervoxels at various scales provide cues for possible groupings of nodes in the CRF to encourage adaptive and long-ranging interactions.
- 330, TITLE: Temporal Action Localization With Pyramid of Score Distribution Features
http://openaccess.thecvf.com/content_cvpr_2016/html/Yuan_Temporal_Action_Localization_CVPR_2016_paper.html
AUTHORS: Jun Yuan, Bingbing Ni, Xiaokang Yang, Ashraf A. Kassim
HIGHLIGHT: We investigate the feature design and classification architectures in temporal action localization.
- 331, TITLE: Recognizing Activities of Daily Living With a Wrist-Mounted Camera
http://openaccess.thecvf.com/content_cvpr_2016/html/Ohnishi_Recognizing_Activities_of_CVPR_2016_paper.html
AUTHORS: Katsunori Ohnishi, Atsushi Kanehira, Asako Kanezaki, Tatsuya Harada
HIGHLIGHT: We present a novel dataset and a novel algorithm for recognizing activities of daily living (ADL) from a first-person wearable camera.
To compare a wrist-mounted camera and a head-mounted camera, we also developed a novel and publicly available dataset that includes videos and annotations of daily activities captured simultaneously by both cameras.
- 332, TITLE: Harnessing Object and Scene Semantics for Large-Scale Video Understanding
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_Harnessing_Object_and_CVPR_2016_paper.html
AUTHORS: Zuxuan Wu, Yanwei Fu, Yu-Gang Jiang, Leonid Sigal
HIGHLIGHT: To address these problems, we propose a novel object- and scene-based semantic fusion network and representation.
- 333, TITLE: Video-Story Composition via Plot Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Choi_Video-Story_Composition_via_CVPR_2016_paper.html
AUTHORS: Jinsoo Choi, Tae-Hyun Oh, In So Kweon
HIGHLIGHT: To efficiently search for the best video-story, we introduce a novel Branch-and-Bound algorithm which guarantees the global optimum.
We collect the dataset consisting of 23 video sets from the web, resulting in a total of 236 individual video clips.
- 334, TITLE: Temporal Action Detection Using a Statistical Language Model
http://openaccess.thecvf.com/content_cvpr_2016/html/Richard_Temporal_Action_Detection_CVPR_2016_paper.html
AUTHORS: Alexander Richard, Juergen Gall
HIGHLIGHT: We propose a novel method for temporal action detection including statistical length and language modeling to represent temporal and contextual structure.
- 335, TITLE: Multi-Scale Patch Aggregation (MPA) for Simultaneous Detection and Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Multi-Scale_Patch_Aggregation_CVPR_2016_paper.html
AUTHORS: Shu Liu, Xiaojuan Qi, Jianping Shi, Hong Zhang, Jiaya Jia
HIGHLIGHT: Aiming at simultaneous detection and segmentation (SDS), we propose a proposal-free framework, which detect and segment object instances via mid-level patches.
- 336, TITLE: Instance-Aware Semantic Segmentation via Multi-Task Network Cascades
http://openaccess.thecvf.com/content_cvpr_2016/html/Dai_Instance-Aware_Semantic_Segmentation_CVPR_2016_paper.html
AUTHORS: Jifeng Dai, Kaiming He, Jian Sun
HIGHLIGHT: In this paper, we present Multi-task Network Cascades for instance-aware semantic segmentation.
- 337, TITLE: ScribbleSup: Scribble-Supervised Convolutional Networks for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Lin_ScribbleSup_Scribble-Supervised_Convolutional_CVPR_2016_paper.html
AUTHORS: Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, Jian Sun

HIGHLIGHT: In this paper, we propose to use scribbles to annotate images, and develop an algorithm to train convolutional networks for semantic segmentation supervised by scribbles.

338, **TITLE:** Feature Space Optimization for Semantic Video Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Kundu_Feature_Space_Optimization_CVPR_2016_paper.html
AUTHORS: Abhijit Kundu, Vibhav Vineet, Vladlen Koltun
HIGHLIGHT: We present an approach to long-range spatio-temporal regularization in semantic video segmentation.

339, **TITLE:** Large-Scale Semantic 3D Reconstruction: An Adaptive Multi-Resolution Model for Multi-Class Volumetric Labeling
http://openaccess.thecvf.com/content_cvpr_2016/html/Blaha_Large-Scale_Semantic_3D_CVPR_2016_paper.html
AUTHORS: Maros Blaha, Christoph Vogel, Audrey Richard, Jan D. Wegner, Thomas Pock, Konrad Schindler
HIGHLIGHT: We propose an adaptive multi-resolution formulation of semantic 3D reconstruction.

340, **TITLE:** Semantic Object Parsing With Local-Global Long Short-Term Memory
http://openaccess.thecvf.com/content_cvpr_2016/html/Liang_Semantic_Object_Parsing_CVPR_2016_paper.html
AUTHORS: Xiaodan Liang, Xiaohui Shen, Donglai Xiang, Jiashi Feng, Liang Lin, Shuicheng Yan
HIGHLIGHT: In this work, we propose a novel deep Local-Global Long Short-Term Memory (LG-LSTM) architecture to seamlessly incorporate short-distance and long-distance spatial dependencies into the feature learning over all pixel positions.

341, **TITLE:** Efficient Piecewise Training of Deep Structured Models for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Lin_Efficient_Piecewise_Training_CVPR_2016_paper.html
AUTHORS: Guosheng Lin, Chunhua Shen, Anton van den Hengel, Ian Reid
HIGHLIGHT: We show how to improve semantic segmentation through the use of contextual information; specifically, we explore 'patch-patch' context between image regions, and 'patch-background' context.

342, **TITLE:** Learning Transferrable Knowledge for Semantic Segmentation With Deep Convolutional Neural Network
http://openaccess.thecvf.com/content_cvpr_2016/html/Hong_Learning_Transferrable_Knowledge_CVPR_2016_paper.html
AUTHORS: Seunghoon Hong, Junhyuk Oh, Honglak Lee, Bohyung Han
HIGHLIGHT: We propose a novel weakly-supervised semantic segmentation algorithm based on Deep Convolutional Neural Network (DCNN).

343, **TITLE:** The Cityscapes Dataset for Semantic Urban Scene Understanding
http://openaccess.thecvf.com/content_cvpr_2016/html/Cordts_The_Cityscapes_Dataset_CVPR_2016_paper.html
AUTHORS: Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, Bero Schiele
HIGHLIGHT: To address this, we introduce Cityscapes, a benchmark suite and large-scale dataset to train and test approaches for pixel-level and instance-level semantic labeling.

344, **TITLE:** Gaussian Conditional Random Field Network for Semantic Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Vemulapalli_Gaussian_Conditional_Random_CVPR_2016_paper.html
AUTHORS: Raviteja Vemulapalli, Oncel Tuzel, Ming-Yu Liu, Rama Chellapa
HIGHLIGHT: In contrast to the existing approaches that use discrete Conditional Random Field (CRF) models, we propose to use a Gaussian CRF model for the task of semantic segmentation.

345, **TITLE:** The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Ros_The_SYNTHIA_Dataset_CVPR_2016_paper.html
AUTHORS: German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, Antonio M. Lopez
HIGHLIGHT: In this paper, we propose to use a virtual world to automatically generate realistic synthetic images with pixel-level annotations.

346, **TITLE:** Progressive Prioritized Multi-View Stereo
http://openaccess.thecvf.com/content_cvpr_2016/html/Locher_Progressive_Prioritized_Multi-View_CVPR_2016_paper.html
AUTHORS: Alex Locher, Michal Perdoch, Luc Van Gool
HIGHLIGHT: This work proposes a progressive patch based multi-view stereo algorithm able to deliver a dense point cloud at any time.

347, **TITLE:** WarpNet: Weakly Supervised Matching for Single-View Reconstruction
http://openaccess.thecvf.com/content_cvpr_2016/html/Kanazawa_WarpNet_Weakly_Supervised_CVPR_2016_paper.html
AUTHORS: Angjoo Kanazawa, David W. Jacobs, Manmohan Chandraker

HIGHLIGHT: We present an approach to matching images of objects in fine-grained datasets without using part annotations, with an application to the challenging problem of weakly supervised single-view reconstruction.

348, **TITLE:** What Sparse Light Field Coding Reveals About Scene Structure
http://openaccess.thecvf.com/content_cvpr_2016/html/Johannsen_What_Sparse_Light_CVPR_2016_paper.html

AUTHORS: Ole Johannsen, Antonin Sulc, Bastian Goldluecke

HIGHLIGHT: In this paper, we propose a novel method for depth estimation in light fields which employs a specifically designed sparse decomposition to leverage the depth-orientation relationship on its epipolar plane images.

349, **TITLE:** Online Reconstruction of Indoor Scenes From RGB-D Streams
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Online_Reconstruction_of_CVPR_2016_paper.html

AUTHORS: Hao Wang, Jun Wang, Wang Liang

HIGHLIGHT: We demonstrate the performance of our system on several real-world scenes and quantitatively assess the modeling accuracy with respect to ground truth models obtained from a LIDAR scanner.

350, **TITLE:** Patches, Planes and Probabilities: A Non-Local Prior for Volumetric 3D Reconstruction
http://openaccess.thecvf.com/content_cvpr_2016/html/Ulusoy_Patches_Planes_and_CVPR_2016_paper.html

AUTHORS: Ali Osman Ulusoy, Michael J. Black, Andreas Geiger

HIGHLIGHT: In this paper, we propose a non-local structured prior for volumetric multi-view 3D reconstruction.

351, **TITLE:** Single Image Camera Calibration With Lenticular Arrays for Augmented Reality
http://openaccess.thecvf.com/content_cvpr_2016/html/Schillebeeckx_Single_Image_Camera_CVPR_2016_paper.html

AUTHORS: Ian Schillebeeckx, Robert Pless

HIGHLIGHT: This paper introduces a calibration objects based on a flat lenticular array that creates a color coded light-field whose observed color changes depending on the angle from which it is viewed.

352, **TITLE:** Augmented Blendshapes for Real-Time Simultaneous 3D Head Modeling and Facial Motion Capture
http://openaccess.thecvf.com/content_cvpr_2016/html/Thomas_Augmented_Blendshapes_for_CVPR_2016_paper.html

AUTHORS: Diego Thomas, Rin-ichiro Taniguchi

HIGHLIGHT: We propose a method to build in real-time animated 3D head models using a consumer-grade RGB-D camera.

353, **TITLE:** Learned Binary Spectral Shape Descriptor for 3D Shape Correspondence
http://openaccess.thecvf.com/content_cvpr_2016/html/Xie_Learned_Binary_Spectral_CVPR_2016_paper.html

AUTHORS: Jin Xie, Meng Wang, Yi Fang

HIGHLIGHT: Different from these real-valued local shape descriptors, in this paper, we propose to learn a novel binary spectral shape descriptor with the deep neural network for 3D shape correspondence.

354, **TITLE:** Multiple Model Fitting as a Set Coverage Problem
http://openaccess.thecvf.com/content_cvpr_2016/html/Magri_Multiple_Model_Fitting_CVPR_2016_paper.html

AUTHORS: Luca Magri, Andrea Fusiello

HIGHLIGHT: We cast the multi-model fitting problem in terms of set covering, deriving a simple and effective method that generalizes Ransac to multiple models and deals with intersecting structures and outliers in a straightforward and principled manner, while avoiding the typical shortcomings of sequential approaches and those of clustering.

355, **TITLE:** Piecewise-Planar 3D Approximation From Wide-Baseline Stereo
http://openaccess.thecvf.com/content_cvpr_2016/html/Verleysen_Piecewise-Planar_3D_Approximation_CVPR_2016_paper.html

AUTHORS: Cedric Verleysen, Christophe De Vleeschouwer

HIGHLIGHT: This paper approximates the 3D geometry of a scene by a small number of 3D planes.

356, **TITLE:** Sparse to Dense 3D Reconstruction From Rolling Shutter Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Saurer_Sparse_to_Dense_CVPR_2016_paper.html

AUTHORS: Olivier Saurer, Marc Pollefeys, Gim Hee Lee

HIGHLIGHT: In this paper, we propose and implement a pipeline for sparse to dense 3D construction with wide baseline images captured from a fast moving rolling shutter camera.

357, **TITLE:** Consistency of Silhouettes and Their Duals
http://openaccess.thecvf.com/content_cvpr_2016/html/Trager_Consistency_of_Silhouettes_CVPR_2016_paper.html

AUTHORS: Matthew Trager, Martial Hebert, Jean Ponce

HIGHLIGHT: In this paper, we investigate the conditions for multiple silhouettes, or more generally arbitrary closed image sets, to be geometrically "consistent".

- 358, TITLE: Rolling Shutter Absolute Pose Problem With Known Vertical Direction
http://openaccess.thecvf.com/content_cvpr_2016/html/Albl_Rolling_Shutter_Absolute_CVPR_2016_paper.html
AUTHORS: Cenek Albl, Zuzana Kukelova, Tomas Pajdla
HIGHLIGHT: We present a solution to the rolling shutter (RS) absolute camera pose problem with known vertical direction.
- 359, TITLE: Uncertainty-Driven 6D Pose Estimation of Objects and Scenes From a Single RGB Image
http://openaccess.thecvf.com/content_cvpr_2016/html/Brachmann_Uncertainty-Driven_6D_Pose_CVPR_2016_paper.html
AUTHORS: Eric Brachmann, Frank Michel, Alexander Krull, Michael Ying Yang, Stefan Gumhold, carsten Rother
HIGHLIGHT: In this work, we show that a single RGB image is sufficient to achieve visually convincing results.
- 360, TITLE: Multicamera Calibration From Visible and Mirrored Epipoles
http://openaccess.thecvf.com/content_cvpr_2016/html/Bushnevskiy_Multicamera_Calibration_From_CVPR_2016_paper.html
AUTHORS: Andrey Bushnevskiy, Lorenzo Sorgi, Bodo Rosenhahn
HIGHLIGHT: In this paper we propose a solution for the automatic recovery of the external calibration of a multicamera system by enforcing only simple geometrical constraints, arising from the epipole visibility, without using any calibration object, such as checkerboards, laser pointers or similar.
- 361, TITLE: Joint Unsupervised Deformable Spatio-Temporal Alignment of Sequences
http://openaccess.thecvf.com/content_cvpr_2016/html/Zafeiriou_Joint_Unsupervised_Deformable_CVPR_2016_paper.html
AUTHORS: Lazaros Zafeiriou, Epameinondas Antonakos, Stefanos Zafeiriou, Maja Pantic
HIGHLIGHT: In this paper, we propose the first, to the best of our knowledge, methodology that can jointly spatio-temporally align two sequences, which display highly deformable texture-varying objects.
- 362, TITLE: Deep Region and Multi-Label Learning for Facial Action Unit Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhao_Deep_Region_and_CVPR_2016_paper.html
AUTHORS: Kaili Zhao, Wen-Sheng Chu, Honggang Zhang
HIGHLIGHT: In this paper, we propose Deep Region and Multi-label Learning (DRML), a unified deep network that simultaneously addresses these two problems.
- 363, TITLE: Constrained Joint Cascade Regression Framework for Simultaneous Facial Action Unit Recognition and Facial Landmark Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_Constrained_Joint_Cascade_CVPR_2016_paper.html
AUTHORS: Yue Wu, Qiang Ji
HIGHLIGHT: In this paper, we improve upon the cascade regression framework and propose the Constrained Joint Cascade Regression Framework (CJCRF) for simultaneous facial action unit recognition and facial landmark detection, which are two related face analysis tasks, but are seldomly exploited together.
- 364, TITLE: Unconstrained Face Alignment via Cascaded Compositional Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Unconstrained_Face_Alignment_CVPR_2016_paper.html
AUTHORS: Shizhan Zhu, Cheng Li, Chen-Change Loy, Xiaoou Tang
HIGHLIGHT: We present a practical approach to address the problem of unconstrained face alignment for a single image.
- 365, TITLE: Automated 3D Face Reconstruction From Multiple Images Using Quality Measures
http://openaccess.thecvf.com/content_cvpr_2016/html/Piotraschke_Automated_3D_Face_CVPR_2016_paper.html
AUTHORS: Marcel Piotraschke, Volker Blanz
HIGHLIGHT: We evaluate different quality measures, develop a method for combining results, and present a complete processing pipeline for automated reconstruction.
- 366, TITLE: Occlusion-Free Face Alignment: Deep Regression Networks Coupled With De-Corrupt AutoEncoders
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Occlusion-Free_Face_Alignment_CVPR_2016_paper.html
AUTHORS: Jie Zhang, Meina Kan, Shiguang Shan, Xilin Chen
HIGHLIGHT: In this work, we propose a novel face alignment method, which cascades several Deep Regression networks coupled with De-corrupt Autoencoders (denoted as DRDA) to explicitly handle partial occlusion problem.
- 367, TITLE: Multimodal Spontaneous Emotion Corpus for Human Behavior Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Multimodal_Spontaneous_Emotion_CVPR_2016_paper.html
AUTHORS: Zheng Zhang, Jeff M. Girard, Yue Wu, Xing Zhang, Peng Liu, Umur Ciftci, Shaun Canavan, Michael Reale, Andy Horowitz, Huiyuan Yang, Jeffrey F. Cohn, Qiang Ji, Lijun Yin
HIGHLIGHT: We present a well-annotated, multimodal, multidimensional spontaneous emotion corpus of 140 participants.

- 368, TITLE: Learning Reconstruction-Based Remote Gaze Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Yu_Learning_Reconstruction-Based_Remote_CVPR_2016_paper.html
AUTHORS: Pei Yu, Jiahuan Zhou, Ying Wu
HIGHLIGHT: A novel method is proposed to learn the metric, such that the affinity structure of the appearance space under this new metric is as close as possible to the affinity structure of the gaze space under the normal Euclidean metric.
- 369, TITLE: Joint Training of Cascaded CNN for Face Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Qin_Joint_Training_of_CVPR_2016_paper.html
AUTHORS: Hongwei Qin, Junjie Yan, Xiu Li, Xiaolin Hu
HIGHLIGHT: In this paper, we propose joint training to achieve end-to-end optimization for CNN cascade.
- 370, TITLE: Facial Expression Intensity Estimation Using Ordinal Information
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhao_Facial_Expression_Intensity_CVPR_2016_paper.html
AUTHORS: Rui Zhao, Quan Gan, Shangfei Wang, Qiang Ji
HIGHLIGHT: In this work, we treat the expression intensity estimation as a regression problem.
- 371, TITLE: Proposal Flow
http://openaccess.thecvf.com/content_cvpr_2016/html/Ham_Proposal_Flow_CVPR_2016_paper.html
AUTHORS: Bumsub Ham, Minsu Cho, Cordelia Schmid, Jean Ponce
HIGHLIGHT: We introduce a novel approach to semantic flow, dubbed proposal flow, that establishes reliable correspondences using object proposals.
We introduce a new dataset that can be used to evaluate both general semantic flow techniques and region-based approaches such as proposal flow.
- 372, TITLE: ProNet: Learning to Propose Object-Specific Boxes for Cascaded Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Sun_ProNet_Learning_to_CVPR_2016_paper.html
AUTHORS: Chen Sun, Manohar Paluri, Ronan Collobert, Ram Nevatia, Lubomir Bourdev
HIGHLIGHT: In this paper, we propose a novel classification architecture ProNet based on convolutional neural networks.
- 373, TITLE: Seeing Behind the Camera: Identifying the Authorship of a Photograph
http://openaccess.thecvf.com/content_cvpr_2016/html/Thomas_Seeing_Behind_the_CVPR_2016_paper.html
AUTHORS: Christopher Thomas, Adriana Kovashka
HIGHLIGHT: We introduce the novel problem of identifying the photographer behind a photograph.
To explore the feasibility of current computer vision techniques to address this problem, we created a new dataset of over 180,000 images taken by 41 well-known photographers.
- 374, TITLE: Material Classification Using Raw Time-Of-Flight Measurements
http://openaccess.thecvf.com/content_cvpr_2016/html/Su_Material_Classification_Using_CVPR_2016_paper.html
AUTHORS: Shuo Chen Su, Felix Heide, Robin Swanson, Jonathan Klein, Clara Callenberg, Matthias Hullin, Wolfgang Heidrich
HIGHLIGHT: We propose a material classification method using raw time-of-flight (ToF) measurements.
- 375, TITLE: Weakly Supervised Object Localization With Progressive Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Weakly_Supervised_Object_CVPR_2016_paper.html
AUTHORS: Dong Li, Jia-Bin Huang, Yali Li, Shengjin Wang, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we address this problem by progressive domain adaptation with two main steps: classification adaptation and detection adaptation.
- 376, TITLE: Newtonian Scene Understanding: Unfolding the Dynamics of Objects in Static Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Mottaghi_Newtonian_Scene_Understanding_CVPR_2016_paper.html
AUTHORS: Roozbeh Mottaghi, Hessam Bagherinezhad, Mohammad Rastegari, Ali Farhadi
HIGHLIGHT: In this paper, we study the challenging problem of predicting the dynamics of objects in static images.
- 377, TITLE: Identifying Good Training Data for Self-Supervised Free Space Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Harakeh_Identifying_Good_Training_CVPR_2016_paper.html
AUTHORS: Ali Harakeh, Daniel Asmar, Elie Shamma
HIGHLIGHT: This paper proposes a novel technique to extract training data from free space in a scene using a stereo camera.

- 378, TITLE: Learning to Match Aerial Images With Deep Attentive Architectures
http://openaccess.thecvf.com/content_cvpr_2016/html/Altwaijry_Learning_to_Match_CVPR_2016_paper.html
AUTHORS: Hani Altwaijry, Eduard Trulls, James Hays, Pascal Fua, Serge Belongie
HIGHLIGHT: In this paper we propose a data-driven, deep learning-based approach that sidesteps local correspondence by framing the problem as a classification task.
- 379, TITLE: Track and Transfer: Watching Videos to Simulate Strong Human Supervision for Weakly-Supervised Object Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Singh_Track_and_Transfer_CVPR_2016_paper.html
AUTHORS: Krishna Kumar Singh, Fanyi Xiao, Yong Jae Lee
HIGHLIGHT: Our framework takes a markedly different direction: we transfer tracked object boxes from weakly-labeled videos to weakly-labeled images to automatically generate pseudo ground-truth boxes, which replace manually annotated bounding boxes.
- 380, TITLE: DeepCAMP: Deep Convolutional Action & Attribute Mid-Level Patterns
http://openaccess.thecvf.com/content_cvpr_2016/html/Diba_DeepCAMP_Deep_Convolutional_CVPR_2016_paper.html
AUTHORS: Ali Diba, Ali Mohammad Pazandeh, Hamed Pirsiavash, Luc Van Gool
HIGHLIGHT: In order to deal with this challenge, we propose a novel convolutional neural network that mines mid-level image patches that are sufficiently dedicated to resolve the corresponding subtleties.
- 381, TITLE: Canny Text Detector: Fast and Robust Scene Text Localization Algorithm
http://openaccess.thecvf.com/content_cvpr_2016/html/Cho_Canny_Text_Detector_CVPR_2016_paper.html
AUTHORS: Hojin Cho, Myungchul Sung, Bongjin Jun
HIGHLIGHT: This paper presents a novel scene text detection algorithm, Canny Text Detector, which takes advantage of the similarity between image edge and text for effective text localization with improved recall rate.
- 382, TITLE: Temporal Multimodal Learning in Audiovisual Speech Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Temporal_Multimodal_Learning_CVPR_2016_paper.html
AUTHORS: Di Hu, Xuelong Li, Xiaoqiang Lu
HIGHLIGHT: In this paper, we will introduce a novel temporal multimodal deep learning architecture, named as Recurrent Temporal Multimodal RBM (RTMRBM), that models multimodal sequences by transforming the sequence of connected MRBMs into a probabilistic series model.
- 383, TITLE: Recovering 6D Object Pose and Predicting Next-Best-View in the Crowd
http://openaccess.thecvf.com/content_cvpr_2016/html/Doumanoglou_Recovering_6D_Object_CVPR_2016_paper.html
AUTHORS: Andreas Doumanoglou, Rigas Kouskouridas, Sotiris Malassiotis, Tae-Kyun Kim
HIGHLIGHT: In this work, we present a complete framework for both single shot-based 6D object pose estimation and next-best-view prediction based on Hough Forests, the state of the art object pose estimator that performs classification and regression jointly.
We provide two additional challenging datasets inspired from realistic scenarios to extensively evaluate the state of the art and our framework.
- 384, TITLE: Robust 3D Hand Pose Estimation in Single Depth Images: From Single-View CNN to Multi-View CNNs
http://openaccess.thecvf.com/content_cvpr_2016/html/Ge_Robust_3D_Hand_CVPR_2016_paper.html
AUTHORS: Lihao Ge, Hui Liang, Junsong Yuan, Daniel Thalmann
HIGHLIGHT: Different from the existing discriminative methods that regress for the hand pose with a single depth image, we propose to first project the query depth image onto three orthogonal planes and utilize these multi-view projections to regress for 2D heat-maps which estimate the joint positions on each plane.
- 385, TITLE: Semantic Segmentation With Boundary Neural Fields
http://openaccess.thecvf.com/content_cvpr_2016/html/Bertasius_Semantic_Segmentation_With_CVPR_2016_paper.html
AUTHORS: Gedas Bertasius, Jianbo Shi, Lorenzo Torresani
HIGHLIGHT: To overcome these problems, we introduce a Boundary Neural Field (BNF), which is a global energy model integrating FCN predictions with boundary cues.
- 386, TITLE: HD Maps: Fine-Grained Road Segmentation by Parsing Ground and Aerial Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Mattyus_HD_Maps_Fine-Grained_CVPR_2016_paper.html
AUTHORS: Gellert Mattyus, Shenlong Wang, Sanja Fidler, Raquel Urtasun
HIGHLIGHT: In this paper we present an approach to enhance existing maps with fine grained segmentation categories such as parking spots and sidewalk, as well as the number and location of road lanes.

- 387, TITLE: DAG-Recurrent Neural Networks For Scene Labeling
http://openaccess.thecvf.com/content_cvpr_2016/html/Shuai_DAG-Recurrent_Neural_Networks_CVPR_2016_paper.html
AUTHORS: Bing Shuai, Zhen Zuo, Bing Wang, Gang Wang
HIGHLIGHT: In this paper, we introduce recurrent neural networks (RNNs) to address this issue.
- 388, TITLE: Saliency Guided Dictionary Learning for Weakly-Supervised Image Parsing
http://openaccess.thecvf.com/content_cvpr_2016/html/Lai_Saliency_Guided_Dictionary_CVPR_2016_paper.html
AUTHORS: Baisheng Lai, Xiaojin Gong
HIGHLIGHT: In this paper, we propose a novel method to perform weakly-supervised image parsing based on the dictionary learning framework.
- 389, TITLE: Attention to Scale: Scale-Aware Semantic Image Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Attention_to_Scale_CVPR_2016_paper.html
AUTHORS: Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, Alan L. Yuille
HIGHLIGHT: In this work, we propose an attention mechanism that learns to softly weight the multi-scale features at each pixel location.
- 390, TITLE: Scene Labeling Using Sparse Precision Matrix
http://openaccess.thecvf.com/content_cvpr_2016/html/Souly_Scene_Labeling_Using_CVPR_2016_paper.html
AUTHORS: Nasim Souly, Mubarak Shah
HIGHLIGHT: In this paper, we propose to use a sparse estimation of precision matrix (also called concentration matrix), which is the inverse of covariance matrix of data obtained by graphical lasso to find interaction between labels and regions.
- 391, TITLE: Iterative Instance Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Iterative_Instance_Segmentation_CVPR_2016_paper.html
AUTHORS: Ke Li, Bharath Hariharan, Jitendra Malik
HIGHLIGHT: We sidestep this problem by reducing structured prediction to a sequence of unconstrained prediction problems and demonstrate that this approach is capable of automatically discovering priors on shape, contiguity of region predictions and smoothness of region contours from data without any a priori specification.
- 392, TITLE: Recurrent Attentional Networks for Saliency Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Kuen_Recurrent_Attentional_Networks_CVPR_2016_paper.html
AUTHORS: Jason Kuen, Zhenhua Wang, Gang Wang
HIGHLIGHT: To overcome such a limitation, in this work, we propose a recurrent attentional convolutional-deconvolution network (RACDNN).
- 393, TITLE: Instance-Level Video Segmentation From Object Tracks
http://openaccess.thecvf.com/content_cvpr_2016/html/Seguin_Instance-Level_Video_Segmentation_CVPR_2016_paper.html
AUTHORS: Guillaume Seguin, Piotr Bojanowski, Remi Lajugie, Ivan Laptev
HIGHLIGHT: We propose a convex relaxation of this problem and solve it efficiently using the Frank-Wolfe algorithm.
- 394, TITLE: Semantic Instance Annotation of Street Scenes by 3D to 2D Label Transfer
http://openaccess.thecvf.com/content_cvpr_2016/html/Xie_Semantic_Instance_Annotation_CVPR_2016_paper.html
AUTHORS: Jun Xie, Martin Kiefel, Ming-Ting Sun, Andreas Geiger
HIGHLIGHT: We start by showing the color coding and label mapping used for the semantic and instance label results in the paper.
- 395, TITLE: Amplitude Modulated Video Camera - Light Separation in Dynamic Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Kolaman_Amplitude_Modulated_Video_CVPR_2016_paper.html
AUTHORS: Amir Kolaman, Maxim Lvov, Rami Hagege, Hugo Guterman
HIGHLIGHT: These condition, caused by a moving shadow for example, force developers to create algorithms which are robust to such variations.
- 396, TITLE: A Benchmark Dataset and Evaluation for Non-Lambertian and Uncalibrated Photometric Stereo
http://openaccess.thecvf.com/content_cvpr_2016/html/Shi_A_Benchmark_Dataset_CVPR_2016_paper.html
AUTHORS: Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, Ping Tan
HIGHLIGHT: Based on our dataset, we quantitatively evaluate state-of-the-art photometric stereo methods for general non-Lambertian materials and unknown lightings to analyze their strengths and limitations.
We then introduce the 'DiLiGenT' photometric stereo image dataset with calibrated Directional Lightings, objects of General reflectance, and 'ground Truth' shapes (normals).

397, TITLE: Depth From Semi-Calibrated Stereo and Defocus
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Depth_From_Semi-Calibrated_CVPR_2016_paper.html
AUTHORS: Ting-Chun Wang, Manohar Srikant, Ravi Ramamoorthi
HIGHLIGHT: In this work, we propose a multi-camera system where we combine a main high-quality camera with two low-res auxiliary cameras.

398, TITLE: Exploiting Spectral-Spatial Correlation for Coded Hyperspectral Image Restoration
http://openaccess.thecvf.com/content_cvpr_2016/html/Fu_Exploiting_Spectral-Spatial_Correlation_CVPR_2016_paper.html
AUTHORS: Ying Fu, Yinqiang Zheng, Imari Sato, Yoichi Sato
HIGHLIGHT: In this paper, we propose an effective method for coded hyperspectral image restoration, which exploits extensive structure sparsity in the hyperspectral image.

399, TITLE: Variable Aperture Light Field Photography: Overcoming the Diffraction-Limited Spatio-Angular Resolution Tradeoff
http://openaccess.thecvf.com/content_cvpr_2016/html/Chang_Variable_Aperture_Light_CVPR_2016_paper.html
AUTHORS: Julie Chang, Isaac Kauvar, Xuemei Hu, Gordon Wetzstein
HIGHLIGHT: In this paper, we analyze fundamental resolution limits of light field cameras in the diffraction limit.

400, TITLE: Convolutional Networks for Shape From Light Field
http://openaccess.thecvf.com/content_cvpr_2016/html/Heber_Convolutional_Networks_for_CVPR_2016_paper.html
AUTHORS: Stefan Heber, Thomas Pock
HIGHLIGHT: In this paper we utilize CNNs to predict depth information for given Light Field (LF) data.

401, TITLE: Panoramic Stereo Videos With a Single Camera
http://openaccess.thecvf.com/content_cvpr_2016/html/Aggarwal_Panoramic_Stereo_Videos_CVPR_2016_paper.html
AUTHORS: Rajat Aggarwal, Amrisha Vohra, Anoop M. Nambodiri
HIGHLIGHT: We present a practical solution for generating 360 degree stereo panoramic videos using a single camera.

402, TITLE: The Next Best Underwater View
http://openaccess.thecvf.com/content_cvpr_2016/html/Sheinin_The_Next_Best_CVPR_2016_paper.html
AUTHORS: Mark Sheinin, Yoav Y. Schechner
HIGHLIGHT: This work proposes to optimize camera and light poses as they move, so that the surface is scanned efficiently and the descattered recovery has the highest quality.

403, TITLE: Reconstructing Shapes and Appearances of Thin Film Objects Using RGB Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Kobayashi_Reconstructing_Shapes_and_CVPR_2016_paper.html
AUTHORS: Yoshie Kobayashi, Tetsuro Morimoto, Imari Sato, Yasuhiro Mukaigawa, Takao Tomono, Katsushi Ikeuchi
HIGHLIGHT: In this paper, we propose a novel method to estimate shapes and film thickness.

404, TITLE: Noisy Label Recovery for Shadow Detection in Unfamiliar Domains
http://openaccess.thecvf.com/content_cvpr_2016/html/Vicente_Noisy_Label_Recovery_CVPR_2016_paper.html
AUTHORS: Tomas F. Yago Vicente, Minh Hoai, Dimitris Samaras
HIGHLIGHT: In this paper we propose "lazy annotation", an efficient annotation method where an annotator only needs to mark the important shadow areas and some non-shadow areas.

405, TITLE: Deep Hand: How to Train a CNN on 1 Million Hand Images When Your Data Is Continuous and Weakly Labelled
http://openaccess.thecvf.com/content_cvpr_2016/html/Koller_Deep_Hand_How_CVPR_2016_paper.html
AUTHORS: Oscar Koller, Hermann Ney, Richard Bowden
HIGHLIGHT: This work presents a new approach to learning a frame-based classifier on weakly labelled sequence data by embedding a CNN within an iterative EM algorithm.

406, TITLE: Recognizing Car Fluents From Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Recognizing_Car_Fluents_CVPR_2016_paper.html
AUTHORS: Bo Li, Tianfu Wu, Caiming Xiong, Song-Chun Zhu
HIGHLIGHT: In this paper, we are interested in inferring the fluents of vehicles from video. Since there are no publicly related dataset, we collect and annotate a car fluent dataset consisting of car videos with diverse fluents.

407, TITLE: Pairwise Decomposition of Image Sequences for Active Multi-View Recognition

- http://openaccess.thecvf.com/content_cvpr_2016/html/Johns_Pairwise_Decomposition_of_CVPR_2016_paper.html
AUTHORS: Edward Johns, Stefan Leutenegger, Andrew J. Davison
HIGHLIGHT: We present state-of-the-art results in both guided and unguided multi-view recognition on the ModelNet dataset, and show how our method can be used with depth images, greyscale images, or both.
- 408, TITLE: Inferring Forces and Learning Human Utilities From Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Inferring_Forces_and_CVPR_2016_paper.html
AUTHORS: Yixin Zhu, Chenfanfu Jiang, Yibiao Zhao, Demetri Terzopoulos, Song-Chun Zhu
HIGHLIGHT: We propose a notion of affordance that takes into account physical quantities generated when the human body interacts with real-world objects, and introduce a learning framework that incorporates the concept of human utilities, which in our opinion provides a deeper and finer-grained account not only of object affordance but also of people's interaction with objects.
- 409, TITLE: Force From Motion: Decoding Physical Sensation in a First Person Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Park_Force_From_Motion_CVPR_2016_paper.html
AUTHORS: Hyun Soo Park, Jyh-Jing Hwang, Jianbo Shi
HIGHLIGHT: In this paper, we focus on a problem of Force from Motion---decoding the sensation of 1) passive forces such as the gravity, 2) the physical scale of the motion (speed) and space, and 3) active forces exerted by the observer such as pedaling a bike or banking on a ski turn.
- 410, TITLE: Robust Multi-Body Feature Tracker: A Segmentation-Free Approach
http://openaccess.thecvf.com/content_cvpr_2016/html/Ji_Robust_Multi-Body_Feature_CVPR_2016_paper.html
AUTHORS: Pan Ji, Hongdong Li, Mathieu Salzmann, Yiran Zhong
HIGHLIGHT: By contrast, here, we introduce a segmentation-free solution to multi-body feature tracking that bypasses the motion assignment step and reduces to solving a series of subproblems with closed-form solutions.
- 411, TITLE: Slow and Steady Feature Analysis: Higher Order Temporal Coherence in Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Jayaraman_Slow_and_Steady_CVPR_2016_paper.html
AUTHORS: Dinesh Jayaraman, Kristen Grauman
HIGHLIGHT: We propose to generalize slow feature analysis to "steady" feature analysis.
- 412, TITLE: Volumetric 3D Tracking by Detection
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Volumetric_3D_Tracking_CVPR_2016_paper.html
AUTHORS: Chun-Hao Huang, Benjamin Allain, Jean-Sebastien Franco, Nassir Navab, Slobodan Ilic, Edmond Boyer
HIGHLIGHT: In this paper, we propose a new framework for 3D tracking by detection based on fully volumetric representations.
- 413, TITLE: The Solution Path Algorithm for Identity-Aware Multi-Object Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Yu_The_Solution_Path_CVPR_2016_paper.html
AUTHORS: Shou-I Yu, Deyu Meng, Wangmeng Zuo, Alexander Hauptmann
HIGHLIGHT: We propose an identity-aware multi-object tracker based on the solution path algorithm.
- 414, TITLE: In Defense of Sparse Tracking: Circulant Sparse Tracker
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_In_Defense_of_CVPR_2016_paper.html
AUTHORS: Tianzhu Zhang, Adel Bibi, Bernard Ghanem
HIGHLIGHT: To deal with the above issues, we propose a novel circulant sparse tracker (CST), which exploits circulant target templates.
- 415, TITLE: Optical Flow With Semantic Segmentation and Localized Layers
http://openaccess.thecvf.com/content_cvpr_2016/html/Sevilla-Lara_Optical_Flow_With_CVPR_2016_paper.html
AUTHORS: Laura Sevilla-Lara, Deqing Sun, Varun Jampani, Michael J. Black
HIGHLIGHT: We define different models of image motion in these regions depending on the type of object. We then pose the flow estimation problem using a novel formulation of localized layers, which addresses limitations of traditional layered models for dealing with complex scene motion.
- 416, TITLE: Video Segmentation via Object Flow
http://openaccess.thecvf.com/content_cvpr_2016/html/Tsai_Video_Segmentation_via_CVPR_2016_paper.html
AUTHORS: Yi-Hsuan Tsai, Ming-Hsuan Yang, Michael J. Black
HIGHLIGHT: To obtain accurate segmentation across time, we propose an efficient algorithm that considers video segmentation and optical flow estimation simultaneously.

- 417, TITLE: Closed-Form Training of Mahalanobis Distance for Supervised Clustering
http://openaccess.thecvf.com/content_cvpr_2016/html/Law_Closed-Form_Training_of_CVPR_2016_paper.html
AUTHORS: Marc T. Law, YaoLiang Yu, Matthieu Cord, Eric P. Xing
HIGHLIGHT: In this paper, we propose a new structured Mahalanobis Distance Metric Learning method for supervised clustering.
- 418, TITLE: Scalable Sparse Subspace Clustering by Orthogonal Matching Pursuit
http://openaccess.thecvf.com/content_cvpr_2016/html/You_Scalable_Sparse_Subspace_CVPR_2016_paper.html
AUTHORS: Chong You, Daniel Robinson, Rene Vidal
HIGHLIGHT: In this paper we study a subspace clustering method based on orthogonal matching pursuit.
- 419, TITLE: Oracle Based Active Set Algorithm for Scalable Elastic Net Subspace Clustering
http://openaccess.thecvf.com/content_cvpr_2016/html/You_Oracle_Based_Active_CVPR_2016_paper.html
AUTHORS: Chong You, Chun-Guang Li, Daniel P. Robinson, Rene Vidal
HIGHLIGHT: This paper studies the geometry of the elastic net regularizer (a mixture of the l_1 and l_2 norms) and uses it to derive a provably correct and scalable active set method for finding the optimal coefficients.
- 420, TITLE: Sparse Coding and Dictionary Learning With Linear Dynamical Systems
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Sparse_Coding_and_CVPR_2016_paper.html
AUTHORS: Wenbing Huang, Fuchun Sun, Lele Cao, Deli Zhao, Huaping Liu, Mehrtash Harandi
HIGHLIGHT: To enhance the performance of LDSs, in this paper, we address the challenging issue of performing sparse coding on the space of LDSs, where both data and dictionary atoms are LDSs.
- 421, TITLE: Sublabel-Accurate Relaxation of Nonconvex Energies
http://openaccess.thecvf.com/content_cvpr_2016/html/Mollenhoff_Sublabel-Accurate_Relaxation_of_CVPR_2016_paper.html
AUTHORS: Thomas Mollenhoff, Emanuel Laude, Michael Moeller, Jan Lellmann, Daniel Cremers
HIGHLIGHT: We propose a novel spatially continuous framework for convex relaxations based on functional lifting.
- 422, TITLE: The Multiverse Loss for Robust Transfer Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Littwin_The_Multiverse_Loss_CVPR_2016_paper.html
AUTHORS: Etai Littwin, Lior Wolf
HIGHLIGHT: In this work we suggest to learn, in the source domain, multiple orthogonal classifiers.
- 423, TITLE: Learning From the Mistakes of Others: Matching Errors in Cross-Dataset Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Sharmanska_Learning_From_the_CVPR_2016_paper.html
AUTHORS: Viktoriia Sharmanska, Novi Quadrianto
HIGHLIGHT: Can we learn about object classes in images by looking at a collection of relevant 3D models?
- 424, TITLE: An Efficient Exact-PGA Algorithm for Constant Curvature Manifolds
http://openaccess.thecvf.com/content_cvpr_2016/html/Chakraborty_An_Efficient_Exact-PGA_CVPR_2016_paper.html
AUTHORS: Rudrasis Chakraborty, Dohyung Seo, Baba C. Vemuri
HIGHLIGHT: In this paper, we propose an efficient exact PGA algorithm for constant curvature Riemannian manifolds (CCM-EPGA).
- 425, TITLE: Online Learning With Bayesian Classification Trees
http://openaccess.thecvf.com/content_cvpr_2016/html/Bulo_Online_Learning_With_CVPR_2016_paper.html
AUTHORS: Samuel Rota Buló, Peter Kotschieder
HIGHLIGHT: In this paper, we propose an online learning algorithm for classification trees that adheres to Bayesian principles.
- 426, TITLE: Cross-Stitch Networks for Multi-Task Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Misra_Cross-Stitch_Networks_for_CVPR_2016_paper.html
AUTHORS: Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, Martial Hebert
HIGHLIGHT: In this paper, we propose a principled approach to learn shared representations in ConvNets using multi-task learning.
- 427, TITLE: Deep Metric Learning via Lifted Structured Feature Embedding
http://openaccess.thecvf.com/content_cvpr_2016/html/Song_Deep_Metric_Learning_CVPR_2016_paper.html
AUTHORS: Hyun Oh Song, Yu Xiang, Stefanie Jegelka, Silvio Savarese

HIGHLIGHT: In this paper, we describe an algorithm for taking full advantage of the training batches in the neural network training by lifting the vector of pairwise distances within the batch to the matrix of pairwise distances. Additionally, we collected Online Products dataset: 120k images of 23k classes of online products for metric learning.

428, **TITLE:** Fast Algorithms for Convolutional Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Lavin_Fast_Algorithms_for_CVPR_2016_paper.html
AUTHORS: Andrew Lavin, Scott Gray
HIGHLIGHT: We introduce a new class of fast algorithms for convolutional neural networks using Winograd's minimal filtering algorithms.

429, **TITLE:** Coordinating Multiple Disparity Proposals for Stereo Computation
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Coordinating_Multiple_Disparity_CVPR_2016_paper.html
AUTHORS: Ang Li, Dapeng Chen, Yuanliu Liu, Zejian Yuan
HIGHLIGHT: In this paper, we alleviate the sensitivity by generating multiple proposals on absolute and relative disparities from multi-segmentations.

430, **TITLE:** Joint Multiview Segmentation and Localization of RGB-D Images Using Depth-Induced Silhouette Consistency
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Joint_Multiview_Segmentation_CVPR_2016_paper.html
AUTHORS: Chi Zhang, Zhiwei Li, Rui Cai, Hongyang Chao, Yong Rui
HIGHLIGHT: In this paper, we propose an RGB-D camera localization approach which takes an effective geometry constraint, i.e. silhouette consistency, into consideration.

431, **TITLE:** A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Mayer_A_Large_Dataset_CVPR_2016_paper.html
AUTHORS: Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, Thomas Brox
HIGHLIGHT: To this end, we propose three synthetic stereo video datasets with sufficient realism, variation, and size to successfully train large networks.

432, **TITLE:** 6D Dynamic Camera Relocalization From Single Reference Image
http://openaccess.thecvf.com/content_cvpr_2016/html/Feng_6D_Dynamic_Camera_CVPR_2016_paper.html
AUTHORS: Wei Feng, Fei-Peng Tian, Qian Zhang, Jizhou Sun
HIGHLIGHT: In this paper, we show that high-quality camera relocalization can be achieved in a much less expensive way.

433, **TITLE:** Dense Monocular Depth Estimation in Complex Dynamic Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Ranftl_Dense_Monocular_Depth_CVPR_2016_paper.html
AUTHORS: Rene Ranftl, Vibhav Vineet, Qifeng Chen, Vladlen Koltun
HIGHLIGHT: We present an approach to dense depth estimation from a single monocular camera that is moving through a dynamic scene.

434, **TITLE:** Using Self-Contradiction to Learn Confidence Measures in Stereo Vision
http://openaccess.thecvf.com/content_cvpr_2016/html/Mostegel_Using_Self-Contradiction_to_CVPR_2016_paper.html
AUTHORS: Christian Mostegel, Markus Rumpel, Friedrich Fraundorfer, Horst Bischof
HIGHLIGHT: To overcome this problem, we propose a new, flexible, and scalable way for generating training data that only requires a set of stereo images as input.

435, **TITLE:** Understanding Real World Indoor Scenes With Synthetic Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Handa_Understanding_Real_World_CVPR_2016_paper.html
AUTHORS: Ankur Handa, Viorica Patraucean, Vijay Badrinarayanan, Simon Stent, Roberto Cipolla
HIGHLIGHT: In this work, we focus our attention on depth based semantic per-pixel labelling as a scene understanding problem and show the potential of computer graphics to generate virtually unlimited labelled data from synthetic 3D scenes.

436, **TITLE:** Stereo Matching With Color and Monochrome Cameras in Low-Light Conditions
http://openaccess.thecvf.com/content_cvpr_2016/html/Jeon_Stereo_Matching_With_CVPR_2016_paper.html
AUTHORS: Hae-Gon Jeon, Joon-Young Lee, Sunghoon Im, Hyowon Ha, In So Kweon
HIGHLIGHT: To address the problem, we present a new stereo matching method with a color and monochrome camera pair.

437, **TITLE:** Camera Calibration From Dynamic Silhouettes Using Motion Barcodes
http://openaccess.thecvf.com/content_cvpr_2016/html/Ben-Artzi_Camera_Calibration_From_CVPR_2016_paper.html
AUTHORS: Gil Ben-Artzi, Yoni Kasten, Shmuel Peleg, Michael Werman

HIGHLIGHT: We propose a speed up of about two orders of magnitude, as well as an increase in robustness and accuracy, to methods computing epipolar geometry from dynamic silhouettes based on a new temporal signature, motion barcode for lines.

438, **TITLE:** Structure-From-Motion Revisited
http://openaccess.thecvf.com/content_cvpr_2016/html/Schonberger_Structure-From-Motion_Revisited_CVPR_2016_paper.html
AUTHORS: Johannes L. Schonberger, Jan-Michael Frahm
HIGHLIGHT: We propose a new SfM technique that improves upon the state of the art to make a further step towards this ultimate goal.

439, **TITLE:** Constructing Canonical Regions for Fast and Effective View Selection
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Constructing_Canonical_Regions_CVPR_2016_paper.html
AUTHORS: Wencheng Wang, Tianhao Gao
HIGHLIGHT: In this paper, we propose a search strategy by identifying the regions that are very likely to contain best views, referred to as canonical regions.

440, **TITLE:** Prior-Less Compressible Structure From Motion
http://openaccess.thecvf.com/content_cvpr_2016/html/Kong_Prior-Less_Compressible_Structure_CVPR_2016_paper.html
AUTHORS: Chen Kong, Simon Lucey
HIGHLIGHT: We argue in this paper that a more expressive and general assumption can be made around compressible 3D structures.

441, **TITLE:** Rolling Shutter Camera Relative Pose: Generalized Epipolar Geometry
http://openaccess.thecvf.com/content_cvpr_2016/html/Dai_Rolling_Shutter_Camera_CVPR_2016_paper.html
AUTHORS: Yuchao Dai, Hongdong Li, Laurent Kneip
HIGHLIGHT: This paper presents a detailed investigation of the geometry of the rolling shutter relative pose problem.

442, **TITLE:** Structure From Motion With Objects
http://openaccess.thecvf.com/content_cvpr_2016/html/Crocco_Structure_From_Motion_CVPR_2016_paper.html
AUTHORS: Marco Crocco, Cosimo Rubino, Alessio Del Bue
HIGHLIGHT: This paper shows for the first time that is possible to reconstruct the position of rigid objects and to jointly recover affine camera calibration solely from a set of object detections in a video sequence.

443, **TITLE:** DeepHand: Robust Hand Pose Estimation by Completing a Matrix Imputed With Deep Features
http://openaccess.thecvf.com/content_cvpr_2016/html/Sinha_DeepHand_Robust_Hand_CVPR_2016_paper.html
AUTHORS: Ayan Sinha, Chiho Choi, Karthik Ramani
HIGHLIGHT: We propose DeepHand to estimate the 3D pose of a hand using depth data from commercial 3D sensors.

444, **TITLE:** Multi-Oriented Text Detection With Fully Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Multi-Oriented_Text_Detection_CVPR_2016_paper.html
AUTHORS: Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, Xiang Bai
HIGHLIGHT: In this paper, we propose an unconventional approach for text detection in natural images.

445, **TITLE:** Robust Scene Text Recognition With Automatic Rectification
http://openaccess.thecvf.com/content_cvpr_2016/html/Shi_Robust_Scene_Text_CVPR_2016_paper.html
AUTHORS: Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, Xiang Bai
HIGHLIGHT: We propose RARE (Robust text recognizer with Automatic RECTification), a recognition model that is robust to irregular text.

446, **TITLE:** Mnemonic Descent Method: A Recurrent Process Applied for End-To-End Face Alignment
http://openaccess.thecvf.com/content_cvpr_2016/html/Trigeorgis_Mnemonic_Descent_Method_CVPR_2016_paper.html
AUTHORS: George Trigeorgis, Patrick Snape, Mihalis A. Nicolaou, Epameinondas Antonakos, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we propose a combined and jointly trained convolutional recurrent neural network architecture that allows the training of an end-to-end to system that attempts to alleviate the aforementioned drawbacks.

447, **TITLE:** Large-Pose Face Alignment via CNN-Based Dense 3D Model Fitting
http://openaccess.thecvf.com/content_cvpr_2016/html/Jourabloo_Large-Pose_Face_Alignment_CVPR_2016_paper.html
AUTHORS: Amin Jourabloo, Xiaoming Liu
HIGHLIGHT: In this paper, we propose a face alignment method for large-pose face images, by combining the powerful cascaded CNN regressor method and 3DMM.

- 448, TITLE: Adaptive 3D Face Reconstruction From Unconstrained Photo Collections
http://openaccess.thecvf.com/content_cvpr_2016/html/Roth_Adaptive_3D_Face_CVPR_2016_paper.html
AUTHORS: Joseph Roth, Yiyong Tong, Xiaoming Liu
HIGHLIGHT: Given a collection of "in-the-wild" face images captured under a variety of unknown pose, expression, and illumination conditions, this paper presents a method for reconstructing a 3D face surface model of an individual along with albedo information.
- 449, TITLE: Online Detection and Classification of Dynamic Hand Gestures With Recurrent 3D Convolutional Neural Network
http://openaccess.thecvf.com/content_cvpr_2016/html/Molchanov_Online_Detection_and_CVPR_2016_paper.html
AUTHORS: Pavlo Molchanov, Xiaodong Yang, Shalini Gupta, Kihwan Kim, Stephen Tyree, Jan Kautz
HIGHLIGHT: In this paper, we address these challenges with a recurrent three-dimensional convolutional neural network that performs simultaneous detection and classification of dynamic hand gestures from multi-modal data. In order to validate our method, we introduce a new challenging multi-modal dynamic hand gesture dataset captured with depth, color and stereo-IR sensors.
- 450, TITLE: Kinematic Structure Correspondences via Hypergraph Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Chang_Kinematic_Structure_Correspondences_CVPR_2016_paper.html
AUTHORS: Hyung Jin Chang, Tobias Fischer, Maxime Petit, Martina Zambelli, Yiannis Demiris
HIGHLIGHT: In this paper, we present a novel framework for finding the kinematic structure correspondence between two objects in videos via hypergraph matching.
- 451, TITLE: CP-mtML: Coupled Projection Multi-Task Metric Learning for Large Scale Face Retrieval
http://openaccess.thecvf.com/content_cvpr_2016/html/Bhattacharai_CP-mtML_Coupled_Projection_CVPR_2016_paper.html
AUTHORS: Binod Bhattacharai, Gaurav Sharma, Frederic Jurie
HIGHLIGHT: We propose a novel Coupled Projection multi-task Metric Learning (CP-mtML) method for large scale face retrieval.
- 452, TITLE: PatchBatch: A Batch Augmented Loss for Optical Flow
http://openaccess.thecvf.com/content_cvpr_2016/html/Gadot_PatchBatch_A_Batch_CVPR_2016_paper.html
AUTHORS: David Gadot, Lior Wolf
HIGHLIGHT: We propose a new pipeline for optical flow computation, based on Deep Learning techniques.
- 453, TITLE: Joint Recovery of Dense Correspondence and Cosegmentation in Two Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Taniai_Joint_Recovery_of_CVPR_2016_paper.html
AUTHORS: Tatsunori Taniai, Sudipta N. Sinha, Yoichi Sato
HIGHLIGHT: We propose a new technique to jointly recover cosegmentation and dense per-pixel correspondence in two images.
- 454, TITLE: Multi-View People Tracking via Hierarchical Trajectory Composition
http://openaccess.thecvf.com/content_cvpr_2016/html/Xu_Multi-View_People_Tracking_CVPR_2016_paper.html
AUTHORS: Yuanlu Xu, Xiaobai Liu, Yang Liu, Song-Chun Zhu
HIGHLIGHT: This paper presents a hierarchical composition approach for multi-view object tracking.
- 455, TITLE: Object Tracking via Dual Linear Structured SVM and Explicit Feature Map
http://openaccess.thecvf.com/content_cvpr_2016/html/Ning_Object_Tracking_via_CVPR_2016_paper.html
AUTHORS: Jifeng Ning, Jimei Yang, Shaojie Jiang, Lei Zhang, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we present a simple yet efficient dual linear SSVM (DLSSVM) algorithm to enable fast learning and execution during tracking.
- 456, TITLE: Robust, Real-Time 3D Tracking of Multiple Objects With Similar Appearances
http://openaccess.thecvf.com/content_cvpr_2016/html/Sekii_Robust_Real-Time_3D_CVPR_2016_paper.html
AUTHORS: Taiki Sekii
HIGHLIGHT: This paper proposes a novel method for tracking multiple moving objects and recovering their three-dimensional (3D) models separately using multiple calibrated cameras.
- 457, TITLE: An Egocentric Look at Video Photographer Identity
http://openaccess.thecvf.com/content_cvpr_2016/html/Hoshen_An_Egocentric_Look_CVPR_2016_paper.html
AUTHORS: Yedid Hoshen, Shmuel Peleg

HIGHLIGHT: An important message in this paper is that photographers should be aware that sharing egocentric video will compromise their anonymity, even when their face is not visible.

458, **TITLE:** Learning Multi-Domain Convolutional Neural Networks for Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Nam_Learning_Multi-Domain_Convolutional_CVPR_2016_paper.html
AUTHORS: Hyeonseob Nam, Bohyung Han
HIGHLIGHT: We propose a novel visual tracking algorithm based on the representations from a discriminatively trained Convolutional Neural Network (CNN).

459, **TITLE:** Hedged Deep Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Qi_Hedged_Deep_Tracking_CVPR_2016_paper.html
AUTHORS: Yuankai Qi, Shengping Zhang, Lei Qin, Hongxun Yao, Qingming Huang, Jongwoo Lim, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a novel CNN based tracking framework, which takes full advantage of features from different CNN layers and uses an adaptive Hedge method to hedge several CNN trackers into a stronger one.

460, **TITLE:** Structural Correlation Filter for Robust Visual Tracking
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Structural_Correlation_Filter_CVPR_2016_paper.html
AUTHORS: Si Liu, Tianzhu Zhang, Xiaochun Cao, Changsheng Xu
HIGHLIGHT: In this paper, we propose a novel structural correlation filter (SCF) model for robust visual tracking.

461, **TITLE:** Visual Tracking Using Attention-Modulated Disintegration and Integration
http://openaccess.thecvf.com/content_cvpr_2016/html/Choi_Visual_Tracking_Using_CVPR_2016_paper.html
AUTHORS: Jongwon Choi, Hyung Jin Chang, Jiyeoup Jeong, Yiannis Demiris, Jin Young Choi
HIGHLIGHT: In this paper, we present a novel attention-modulated visual tracking algorithm that decomposes an object into multiple cognitive units, and trains multiple elementary trackers in order to modulate the distribution of attention according to various feature and kernel types.

462, **TITLE:** A Continuous Occlusion Model for Road Scene Understanding
http://openaccess.thecvf.com/content_cvpr_2016/html/Dhiman_A_Continuous_Occlusion_CVPR_2016_paper.html
AUTHORS: Vikas Dhiman, Quoc-Huy Tran, Jason J. Corso, Manmohan Chandraker
HIGHLIGHT: We present a physically interpretable, continuous 3D model for handling occlusions with applications to road scene understanding.

463, **TITLE:** Virtual Worlds as Proxy for Multi-Object Tracking Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Gaidon_Virtual_Worlds_as_CVPR_2016_paper.html
AUTHORS: Adrien Gaidon, Qiao Wang, Yohann Cabon, Eleonora Vig
HIGHLIGHT: We propose an efficient real-to-virtual world cloning method, and validate our approach by building and publicly releasing a new video dataset, called Virtual KITTI, automatically labeled with accurate ground truth for object detection, tracking, scene and instance segmentation, depth, and optical flow.

464, **TITLE:** Uncalibrated Photometric Stereo by Stepwise Optimization Using Principal Components of Isotropic BRDFs
http://openaccess.thecvf.com/content_cvpr_2016/html/Midorikawa_Uncalibrated_Photometric_Stereo_CVPR_2016_paper.html
AUTHORS: Keisuke Midorikawa, Toshihiko Yamasaki, Kiyoharu Aizawa
HIGHLIGHT: We propose a model that represents various isotropic reflectance functions by using the principal components of items in a dataset, and formulate the uncalibrated photometric stereo as a regression problem.

465, **TITLE:** Unbiased Photometric Stereo for Colored Surfaces: A Variational Approach
http://openaccess.thecvf.com/content_cvpr_2016/html/Queau_Unbiased_Photometric_Stereo_CVPR_2016_paper.html
AUTHORS: Yvain Queau, Roberto Mecca, Jean-Denis Durou
HIGHLIGHT: In this work we propose a new formulation of color photometric stereo, based on image ratios, that makes the technique independent from the albedos.

466, **TITLE:** 3D Reconstruction of Transparent Objects With Position-Normal Consistency
http://openaccess.thecvf.com/content_cvpr_2016/html/Qian_3D_Reconstruction_of_CVPR_2016_paper.html
AUTHORS: Yiming Qian, Minglun Gong, Yee Hong Yang
HIGHLIGHT: Under the assumption that the rays refract only twice when traveling through the object, we present the first approach to simultaneously reconstructing the 3D positions and normals of the object's surface at both refraction locations.

467, **TITLE:** Real-Time Depth Refinement for Specular Objects
http://openaccess.thecvf.com/content_cvpr_2016/html/Or-El_Real-Time_Depth_Refinement_CVPR_2016_paper.html

AUTHORS: Roy Or-El, Rom Hershkovitz, Aaron Wetzler, Guy Rosman, Alfred M. Bruckstein, Ron Kimmel
HIGHLIGHT: We take advantage of the built-in monochromatic IR projector and IR images of the RGB-D scanners and present a lighting model that accounts for the specular regions in the input image.

468, TITLE: Recovering Transparent Shape From Time-Of-Flight Distortion
http://openaccess.thecvf.com/content_cvpr_2016/html/Tanaka_Recovering_Transparent_Shape_CVPR_2016_paper.html
AUTHORS: Kenichiro Tanaka, Yasuhiro Mukaigawa, Hiroyuki Kubo, Yasuyuki Matsushita, Yasushi Yagi
HIGHLIGHT: This paper presents a method for recovering shape and normal of a transparent object from a single viewpoint using a Time-of-Flight (ToF) camera.

469, TITLE: Robust Light Field Depth Estimation for Noisy Scene With Occlusion
http://openaccess.thecvf.com/content_cvpr_2016/html/Williem_Robust_Light_Field_CVPR_2016_paper.html
AUTHORS: W. Williem, In Kyu Park
HIGHLIGHT: To remedy this problem, we present a light field depth estimation method which is more robust to occlusion and less sensitive to noise.

470, TITLE: Rotational Crossed-Slit Light Field
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Rotational_Crossed-Slit_Light_CVPR_2016_paper.html
AUTHORS: Nianyi Li, Haiting Lin, Bilin Sun, Mingyuan Zhou, Jingyi Yu
HIGHLIGHT: In this paper, we present a novel LF sampling scheme by exploiting a special non-centric camera called the crossed-slit or XSlit camera.

471, TITLE: Single Image Object Modeling Based on BRDF and R-Surfaces Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Natola_Single_Image_Object_CVPR_2016_paper.html
AUTHORS: Fabrizio Natola, Valsamis Ntouskos, Fiora Pirri, Marta Sanzari
HIGHLIGHT: Single Image Object Modeling Based on BRDF and R-Surfaces Learning

472, TITLE: A Nonlinear Regression Technique for Manifold Valued Data With Applications to Medical Image Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Banerjee_A_Nonlinear_Regression_CVPR_2016_paper.html
AUTHORS: Monami Banerjee, Rudrasis Chakraborty, Edward Ofori, Michael S. Okun, David E. Viallancourt, Baba C. Vemuri
HIGHLIGHT: In this paper, we present a novel nonlinear kernel-based regression method that is applicable to manifold valued data.
We present an extensive set of results along with statistical validation and comparisons.

473, TITLE: RAID-G: Robust Estimation of Approximate Infinite Dimensional Gaussian With Application to Material Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_RAID-G_Robust_Estimation_CVPR_2016_paper.html
AUTHORS: Qilong Wang, Peihua Li, Wangmeng Zuo, Lei Zhang
HIGHLIGHT: In this paper, we propose a novel image descriptor, namely, robust approximate infinite dimensional Gaussian (RAID-G).

474, TITLE: An Empirical Evaluation of Current Convolutional Architectures' Ability to Manage Nuisance Location and Scale Variability
http://openaccess.thecvf.com/content_cvpr_2016/html/Karianakis_An_Empirical_Evaluation_CVPR_2016_paper.html
AUTHORS: Nikolaos Karianakis, Jingming Dong, Stefano Soatto
HIGHLIGHT: We also quantify the effects of context on the overall classification task and its impact on the performance of CNNs, and propose improved sampling techniques for heuristic proposal schemes that improve end-to-end performance to state-of-the-art levels.

475, TITLE: Learning Sparse High Dimensional Filters: Image Filtering, Dense CRFs and Bilateral Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Jampani_Learning_Sparse_High_CVPR_2016_paper.html
AUTHORS: Varun Jampani, Martin Kiefel, Peter V. Gehler
HIGHLIGHT: In this paper, we will generalize the parametrization and in particular derive a gradient descent algorithm so the filter parameters can be learned from data.

476, TITLE: Mixture of Bilateral-Projection Two-Dimensional Probabilistic Principal Component Analysis
http://openaccess.thecvf.com/content_cvpr_2016/html/Ju_Mixture_of_Bilateral-Projection_CVPR_2016_paper.html
AUTHORS: Fujiao Ju, Yanfeng Sun, Junbin Gao, Simeng Liu, Yongli Hu, Baocai Yin
HIGHLIGHT: This paper proposes a mixture of bilateral-projection probabilistic principal component analysis model (mixB2DPPCA) on 2D data.

- 477, TITLE: Rolling Rotations for Recognizing Human Actions From 3D Skeletal Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Vemulapalli_Rolling_Rotations_for_CVPR_2016_paper.html
AUTHORS: Raviteja Vemulapalli, Rama Chellapa
HIGHLIGHT: In this work, we use rolling maps for recognizing human actions from 3D skeletal data.
- 478, TITLE: Improving the Robustness of Deep Neural Networks via Stability Training
http://openaccess.thecvf.com/content_cvpr_2016/html/Zheng_Improving_the_Robustness_CVPR_2016_paper.html
AUTHORS: Stephan Zheng, Yang Song, Thomas Leung, Ian Goodfellow
HIGHLIGHT: In this paper we address the issue of output instability of deep neural networks: small perturbations in the visual input can significantly distort the feature embeddings and output of a neural network.
- 479, TITLE: Logistic Boosting Regression for Label Distribution Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Xing_Logistic_Boosting_Regression_CVPR_2016_paper.html
AUTHORS: Chao Xing, Xin Geng, Hui Xue
HIGHLIGHT: The base learners are chosen as weighted regression tree and vector tree, which constitute two algorithms named LDLogitBoost and AOSO-LDLogitBoost in this paper.
- 480, TITLE: Efficient Temporal Sequence Comparison and Classification Using Gram Matrix Embeddings on a Riemannian Manifold
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Efficient_Temporal_Sequence_CVPR_2016_paper.html
AUTHORS: Xikang Zhang, Yin Wang, Mengran Gou, Mario Sznaier, Octavia Camps
HIGHLIGHT: In this paper we propose a new framework to compare and classify temporal sequences.
- 481, TITLE: Deep Reflectance Maps
http://openaccess.thecvf.com/content_cvpr_2016/html/Rematas_Deep_Reflectance_Maps_CVPR_2016_paper.html
AUTHORS: Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Efstratios Gavves, Tinne Tuytelaars
HIGHLIGHT: We propose a fully convolutional neural architecture to estimate reflectance maps of specular materials in natural lighting conditions.
In order to analyze performance on this difficult task, we propose a new challenge of Specular MATerials on SHapes with complex IllumiNation (SMASHINg) using both synthetic and real images.
- 482, TITLE: Semantic Filtering
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Semantic_Filtering_CVPR_2016_paper.html
AUTHORS: Qingxiong Yang
HIGHLIGHT: Inspired by the fact that learning-based edge detectors/classifiers significantly outperform traditional manually-designed detectors, this paper proposes a learning-based edge-preserving filtering technique.
- 483, TITLE: UAV Sensor Fusion With Latent-Dynamic Conditional Random Fields in Coronal Plane Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Rahimi_UAV_Sensor_Fusion_CVPR_2016_paper.html
AUTHORS: Amir M. Rahimi, Raphael Ruschel, B.S. Manjunath
HIGHLIGHT: We present a real-time body orientation estimation in a micro-Unmanned Air Vehicle video stream.
- 484, TITLE: Robust Visual Place Recognition With Graph Kernels
http://openaccess.thecvf.com/content_cvpr_2016/html/Stumm_Robust_Visual_Place_CVPR_2016_paper.html
AUTHORS: Elena Stumm, Christopher Mei, Simon Lacroix, Juan Nieto, Marco Hutter, Roland Siegwart
HIGHLIGHT: A novel method for visual place recognition is introduced and evaluated, demonstrating robustness to perceptual aliasing and observation noise.
- 485, TITLE: Semantic Image Segmentation With Task-Specific Edge Detection Using CNNs and a Discriminatively Trained Domain Transform
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Semantic_Image_Segmentation_CVPR_2016_paper.html
AUTHORS: Liang-Chieh Chen, Jonathan T. Barron, George Papandreou, Kevin Murphy, Alan L. Yuille
HIGHLIGHT: We propose replacing the fully-connected CRF with domain transform (DT), a modern edge-preserving filtering method in which the amount of smoothing is controlled by a reference edge map.
- 486, TITLE: Natural Language Object Retrieval
http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Natural_Language_Object_CVPR_2016_paper.html
AUTHORS: Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, Trevor Darrell

HIGHLIGHT: In this paper, we address the task of natural language object retrieval, to localize a target object within a given image based on a natural language query of the object.

487, **TITLE:** DenseCap: Fully Convolutional Localization Networks for Dense Captioning
http://openaccess.thecvf.com/content_cvpr_2016/html/Johnson_DenseCap_Fully_Convolutional_CVPR_2016_paper.html
AUTHORS: Justin Johnson, Andrej Karpathy, Li Fei-Fei
HIGHLIGHT: We introduce the dense captioning task, which requires a computer vision system to both localize and describe salient regions in images in natural language.

488, **TITLE:** Unsupervised Learning From Narrated Instruction Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Alayrac_Unsupervised_Learning_From_CVPR_2016_paper.html
AUTHORS: Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, Simon Lacoste-Julien
HIGHLIGHT: The contributions of this paper are three-fold.
Second, we collect and annotate a new challenging dataset of real-world instruction videos from the Internet.

489, **TITLE:** Video Paragraph Captioning Using Hierarchical Recurrent Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Yu_Video_Paragraph_Captioning_CVPR_2016_paper.html
AUTHORS: Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, Wei Xu
HIGHLIGHT: We present an approach that exploits hierarchical Recurrent Neural Networks (RNNs) to tackle the video captioning problem, i.e., generating one or multiple sentences to describe a realistic video.

490, **TITLE:** Jointly Modeling Embedding and Translation to Bridge Video and Language
http://openaccess.thecvf.com/content_cvpr_2016/html/Pan_Jointly_Modeling_Embedding_CVPR_2016_paper.html
AUTHORS: Yingwei Pan, Tao Mei, Ting Yao, Houqiang Li, Yong Rui
HIGHLIGHT: This paper presents a novel unified framework, named Long Short-Term Memory with visual-semantic Embedding (LSTM-E), which can simultaneously explore the learning of LSTM and visual-semantic embedding.

491, **TITLE:** We Are Humor Beings: Understanding and Predicting Visual Humor
http://openaccess.thecvf.com/content_cvpr_2016/html/Chandrasekaran_We_Are_Humor_CVPR_2016_paper.html
AUTHORS: Arjun Chandrasekaran, Ashwin K. Vijayakumar, Stanislaw Antol, Mohit Bansal, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh
HIGHLIGHT: In this work, we are interested in the question - what content in a scene causes it to be funny?
We collect two datasets of abstract scenes that facilitate the study of humor at both the scene-level and the object-level.

492, **TITLE:** Where to Look: Focus Regions for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2016/html/Shih_Where_to_Look_CVPR_2016_paper.html
AUTHORS: Kevin J. Shih, Saurabh Singh, Derek Hoiem
HIGHLIGHT: We present a method that learns to answer visual questions by selecting image regions relevant to the text-based query.

493, **TITLE:** Ask Me Anything: Free-Form Visual Question Answering Based on Knowledge From External Sources
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_Ask_MeAnything_CVPR_2016_paper.html
AUTHORS: Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, Anton van den Hengel
HIGHLIGHT: We propose a method for visual question answering which combines an internal representation of the content of an image with information extracted from a general knowledge base to answer a broad range of image-based questions.

494, **TITLE:** MovieQA: Understanding Stories in Movies Through Question-Answering
http://openaccess.thecvf.com/content_cvpr_2016/html/Tapaswi_MovieQA_Understanding_Stories_CVPR_2016_paper.html
AUTHORS: Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, Sanja Fidler
HIGHLIGHT: We introduce the MovieQA dataset which aims to evaluate automatic story comprehension from both video and text.

495, **TITLE:** TGIF: A New Dataset and Benchmark on Animated GIF Description
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_TGIF_A_New_CVPR_2016_paper.html
AUTHORS: Yuncheng Li, Yale Song, Liangliang Cao, Joel Tetreault, Larry Goldberg, Alejandro Jaimes, Jiebo Luo
HIGHLIGHT: We perform extensive statistical analyses to compare our dataset to existing image and video description datasets.
To advance research on animated GIF understanding, we collected a new dataset, Tumblr GIF (TGIF), with 100K animated GIFs from Tumblr and 120K natural language descriptions obtained via crowdsourcing.

- 496, TITLE: Image Captioning With Semantic Attention
http://openaccess.thecvf.com/content_cvpr_2016/html/You_Image_Captioning_With_CVPR_2016_paper.html
AUTHORS: Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, Jiebo Luo
HIGHLIGHT: In this paper, we propose a new algorithm that combines both approaches through a model of semantic attention.
- 497, TITLE: Temporally Coherent 4D Reconstruction of Complex Dynamic Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Mustafa_Temporally_Coherent_4D_CVPR_2016_paper.html
AUTHORS: Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, Adrian Hilton
HIGHLIGHT: This paper presents an approach for reconstruction of 4D temporally coherent models of complex dynamic scenes.
- 498, TITLE: Consensus of Non-Rigid Reconstructions
http://openaccess.thecvf.com/content_cvpr_2016/html/Lee_Consensus_of_Non-Rigid_CVPR_2016_paper.html
AUTHORS: Minsik Lee, Jungchan Cho, Songhwai Oh
HIGHLIGHT: In order to resolve the reflection ambiguity between weak (and possibly bad) reconstructions, we propose a novel optimization framework which only involves a single eigenvalue decomposition.
- 499, TITLE: Isometric Non-Rigid Shape-From-Motion in Linear Time
http://openaccess.thecvf.com/content_cvpr_2016/html/Parashar_Isometric_Non-Rigid_Shape-From-Motion_CVPR_2016_paper.html
AUTHORS: Shaifali Parashar, Daniel Pizarro, Adrien Bartoli
HIGHLIGHT: We propose a new theoretical framework based on Riemmanian manifolds to represent the unknown 3D surfaces, as embeddings of the camera's retinal planes.
- 500, TITLE: Learning Online Smooth Predictors for Realtime Camera Planning Using Recurrent Decision Trees
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Learning_Online_Smooth_CVPR_2016_paper.html
AUTHORS: Jianhui Chen, Hoang M. Le, Peter Carr, Yisong Yue, James J. Little
HIGHLIGHT: We propose a recurrent decision tree framework that can directly incorporate temporal consistency into a data-driven predictor, as well as a learning algorithm that can efficiently learn such temporally smooth models.
- 501, TITLE: Egocentric Future Localization
http://openaccess.thecvf.com/content_cvpr_2016/html/Park_Egocentric_Future_Localization_CVPR_2016_paper.html
AUTHORS: Hyun Soo Park, Jyh-Jing Hwang, Yedong Niu, Jianbo Shi
HIGHLIGHT: We presents a method for future localization: to predict plausible future trajectories of ego-motion in egocentric stereo images.
- 502, TITLE: Full Flow: Optical Flow Estimation By Global Optimization Over Regular Grids
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Full_Flow_Optical_CVPR_2016_paper.html
AUTHORS: Qifeng Chen, Vladlen Koltun
HIGHLIGHT: We present a global optimization approach to optical flow estimation.
- 503, TITLE: Structured Feature Learning for Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Chu_Structured_Feature_Learning_CVPR_2016_paper.html
AUTHORS: Xiao Chu, Wanli Ouyang, Hongsheng Li, Xiaogang Wang
HIGHLIGHT: In this paper, we propose a structured feature learning framework to reason the correlation among body joints at the feature level in human pose estimation.
- 504, TITLE: Convolutional Pose Machines
http://openaccess.thecvf.com/content_cvpr_2016/html/Wei_Convolutional_Pose_Machines_CVPR_2016_paper.html
AUTHORS: Shih-En Wei, Varun Ramakrishna, Takeo Kanade, Yaser Sheikh
HIGHLIGHT: In this work we show a systematic design for how convolutional networks can be incorporated into the pose machine framework for learning image features and image-dependent spatial models for the task of pose estimation.
- 505, TITLE: Human Pose Estimation With Iterative Error Feedback
http://openaccess.thecvf.com/content_cvpr_2016/html/Carreira_Human_Pose_Estimation_CVPR_2016_paper.html
AUTHORS: Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, Jitendra Malik
HIGHLIGHT: Here we propose a framework that expands the expressive power of hierarchical feature extractors to encompass both input and output spaces, by introducing top-down feedback.
- 506, TITLE: WELDON: Weakly Supervised Learning of Deep Convolutional Neural Networks

http://openaccess.thecvf.com/content_cvpr_2016/html/Durand_WELDON_Weakly_Supervised_CVPR_2016_paper.html
AUTHORS: Thibaut Durand, Nicolas Thome, Matthieu Cord
HIGHLIGHT: In this paper, we introduce a novel framework for WEakly supervised Learning of Deep cONvolutional neural Networks (WELDON).

507, TITLE: DisturbLabel: Regularizing CNN on the Loss Layer
http://openaccess.thecvf.com/content_cvpr_2016/html/Xie_DisturbLabel_Regularizing_CNN_CVPR_2016_paper.html
AUTHORS: Lingxi Xie, Jingdong Wang, Zhen Wei, Meng Wang, Qi Tian
HIGHLIGHT: In this paper, we present DisturbLabel, an extremely simple algorithm which randomly replaces a part of labels as incorrect values in each iteration.

508, TITLE: Gradual DropIn of Layers to Train Very Deep Neural Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Smith_Gradual_DropIn_of_CVPR_2016_paper.html
AUTHORS: Leslie N. Smith, Emily M. Hand, Timothy Doster
HIGHLIGHT: We introduce the concept of dynamically growing a neural network during training.

509, TITLE: Structure Inference Machines: Recurrent Neural Networks for Analyzing Relations in Group Activity Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Deng_Structure_Inference_Machines_CVPR_2016_paper.html
AUTHORS: Zhiwei Deng, Arash Vahdat, Hexiang Hu, Greg Mori
HIGHLIGHT: In this paper, we propose a method to integrate graphical models and deep neural networks into a joint framework.

510, TITLE: Deep SimNets
http://openaccess.thecvf.com/content_cvpr_2016/html/Cohen_Deep_SimNets_CVPR_2016_paper.html
AUTHORS: Nadav Cohen, Or Sharir, Amnon Shashua
HIGHLIGHT: We present a deep layered architecture that generalizes convolutional neural networks (ConvNets).

511, TITLE: Studying Very Low Resolution Recognition Using Deep Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Studying_Very_Low_CVPR_2016_paper.html
AUTHORS: Zhangyang Wang, Shiyu Chang, Yingzhen Yang, Ding Liu, Thomas S. Huang
HIGHLIGHT: Taking advantage of techniques primarily in super resolution, domain adaptation and robust regression, we formulate a dedicated deep learning method and demonstrate how these techniques are incorporated step by step.

512, TITLE: Deep Gaussian Conditional Random Field Network: A Model-Based Deep Network for Discriminative Denoising
http://openaccess.thecvf.com/content_cvpr_2016/html/Vemulapalli_Deep_Gaussian_Conditional_CVPR_2016_paper.html
AUTHORS: Raviteja Vemulapalli, Oncel Tuzel, Ming-Yu Liu
HIGHLIGHT: We propose a novel end-to-end trainable deep network architecture for image denoising based on a Gaussian Conditional Random Field (GCRF) model.

513, TITLE: Event-Specific Image Importance
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Event-Specific_Image_Importance_CVPR_2016_paper.html
AUTHORS: Yufei Wang, Zhe Lin, Xiaohui Shen, Radomir Mech, Gavin Miller, Garrison W. Cottrell
HIGHLIGHT: In this paper, we show that the selection of important images is consistent among different viewers, and that this selection process is related to the event type of the album.
We collected a new event album dataset with human annotation of the relative image importance with each event album.

514, TITLE: Quantized Convolutional Neural Networks for Mobile Devices
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_Quantized_Convolutional_Neural_CVPR_2016_paper.html
AUTHORS: Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, Jian Cheng
HIGHLIGHT: In this paper, we propose an efficient framework, namely Quantized CNN, to simultaneously speed-up the computation and reduce the storage and memory overhead of CNN models.

515, TITLE: Inverting Visual Representations With Convolutional Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Dosovitskiy_Inverting_Visual_Representations_CVPR_2016_paper.html
AUTHORS: Alexey Dosovitskiy, Thomas Brox
HIGHLIGHT: We propose a new approach to study image representations by inverting them with an up-convolutional neural network.

- 516, TITLE: Pose-Aware Face Recognition in the Wild
http://openaccess.thecvf.com/content_cvpr_2016/html/Masi_Pose-Aware_Face_Recognition_CVPR_2016_paper.html
AUTHORS: Iacopo Masi, Stephen Rawls, Gerard Medioni, Prem Natarajan
HIGHLIGHT: We propose a method to push the frontiers of unconstrained face recognition in the wild, focusing on the problem of extreme pose variations.
- 517, TITLE: Multi-View Deep Network for Cross-View Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Kan_Multi-View_Deep_Network_CVPR_2016_paper.html
AUTHORS: Meina Kan, Shiguang Shan, Xilin Chen
HIGHLIGHT: To eliminate the complex (maybe even highly nonlinear) view discrepancy for favorable cross-view recognition, we propose a multi-view deep network (MvDN), which seeks for a non-linear discriminant and view-invariant representation shared between multiple views.
- 518, TITLE: Sparsifying Neural Network Connections for Face Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Sun_Sparsifying_Neural_Network_CVPR_2016_paper.html
AUTHORS: Yi Sun, Xiaogang Wang, Xiaoou Tang
HIGHLIGHT: This paper proposes to learn high-performance deep ConvNets with sparse neural connections, referred to as sparse ConvNets, for face recognition.
- 519, TITLE: Pairwise Linear Regression Classification for Image Set Retrieval
http://openaccess.thecvf.com/content_cvpr_2016/html/Feng_Pairwise_Linear_Regression_CVPR_2016_paper.html
AUTHORS: Qingxiang Feng, Yicong Zhou, Rushi Lan
HIGHLIGHT: This paper proposes the pairwise linear regression classification (PLRC) for image set retrieval.
- 520, TITLE: The MegaFace Benchmark: 1 Million Faces for Recognition at Scale
http://openaccess.thecvf.com/content_cvpr_2016/html/Kemelmacher-Shlizerman_The_MegaFace_Benchmark_CVPR_2016_paper.html
AUTHORS: Ira Kemelmacher-Shlizerman, Steven M. Seitz, Daniel Miller, Evan Brossard
HIGHLIGHT: In this paper, we advocate evaluations at the million scale (LFW includes only 13K photos of 5K people). To this end, we have assembled the MegaFace dataset and created the first MegaFace challenge.
- 521, TITLE: Learnt Quasi-Transitive Similarity for Retrieval From Large Collections of Faces
http://openaccess.thecvf.com/content_cvpr_2016/html/Arandjelovic_Learnt_Quasi-Transitive_Similarity_CVPR_2016_paper.html
AUTHORS: Ognjen Arandjelovic
HIGHLIGHT: Given a baseline algorithm for measuring the similarity of two face sets, the meta-algorithm introduced in this paper seeks to leverage the structure of the data corpus to make the best use of the available baseline.
- 522, TITLE: Latent Factor Guided Convolutional Neural Networks for Age-Invariant Face Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Wen_Latent_Factor_Guided_CVPR_2016_paper.html
AUTHORS: Yandong Wen, Zhifeng Li, Yu Qiao
HIGHLIGHT: In order to address this problem, we propose a novel deep face recognition framework to learn the age-invariant deep face features through a carefully designed CNN model.
- 523, TITLE: Copula Ordinal Regression for Joint Estimation of Facial Action Unit Intensity
http://openaccess.thecvf.com/content_cvpr_2016/html/Walecki_Copula_Ordinal_Regression_CVPR_2016_paper.html
AUTHORS: Robert Walecki, Ognjen Rudovic, Vladimir Pavlovic, Maja Pantic
HIGHLIGHT: To this end, we introduce a novel modeling framework, Copula Ordinal Regression (COR), that leverages the power of copula functions and CRFs, to detangle the probabilistic modeling of AU dependencies from the marginal modeling of the AU intensity.
- 524, TITLE: A Robust Multilinear Model Learning Framework for 3D Faces
http://openaccess.thecvf.com/content_cvpr_2016/html/Bolkart_A_Robust_Multilinear_CVPR_2016_paper.html
AUTHORS: Timo Bolkart, Stefanie Wuhrer
HIGHLIGHT: To overcome this, we introduce the first framework to robustly learn a multilinear model from 3D face databases with missing data, corrupt data, wrong semantic correspondence, and inaccurate vertex correspondence.
- 525, TITLE: Ordinal Regression With Multiple Output CNN for Age Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Niu_Ordinal_Regression_With_CVPR_2016_paper.html
AUTHORS: Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, Gang Hua
HIGHLIGHT: In this paper, we propose an End-to-End learning approach to address ordinal regression problems using deep Convolutional Neural Network, which could simultaneously conduct feature learning and regression modeling.

- 526, TITLE: DeepCut: Joint Subset Partition and Labeling for Multi Person Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Pishchulin_DeepCut_Joint_Subset_CVPR_2016_paper.html
AUTHORS: Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Bjoern Andres, Mykhaylo Andriluka, Peter V. Gehler, Bernt Schiele
HIGHLIGHT: We propose an approach that jointly solves the tasks of detection and pose estimation: it infers the number of persons in a scene, identifies occluded body parts, and disambiguates body parts between people in close proximity of each other.
- 527, TITLE: Thin-Slicing for Pose: Learning to Understand Pose Without Explicit Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Kwak_Thin-Slicing_for_Pose_CVPR_2016_paper.html
AUTHORS: Suha Kwak, Minsu Cho, Ivan Laptev
HIGHLIGHT: We address the problem of learning a pose-aware, compact embedding that projects images with similar human poses to be placed close-by in the embedding space.
- 528, TITLE: A Dual-Source Approach for 3D Pose Estimation From a Single Image
http://openaccess.thecvf.com/content_cvpr_2016/html/Yasin_A_Dual-Source_Approach_CVPR_2016_paper.html
AUTHORS: Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, Juergen Gall
HIGHLIGHT: We therefore propose to use two independent training sources.
- 529, TITLE: Efficiently Creating 3D Training Data for Fine Hand Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Oberweger_Efficiently_Creating_3D_CVPR_2016_paper.html
AUTHORS: Markus Oberweger, Gernot Riegler, Paul Wohlhart, Vincent Lepetit
HIGHLIGHT: We propose a semi-automated method for efficiently and accurately labeling each frame of a hand depth video with the corresponding 3D locations of the joints: The user is asked to provide only an estimate of the 2D reprojections of the visible joints in some reference frames, which are automatically selected to minimize the labeling work by efficiently optimizing a sub-modular loss function.
- 530, TITLE: Sparseness Meets Deepness: 3D Human Pose Estimation From Monocular Video
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Sparseness_Meets_Deepness_CVPR_2016_paper.html
AUTHORS: Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G. Derpanis, Kostas Daniilidis
HIGHLIGHT: In the former case, a novel approach is introduced that integrates a sparsity-driven 3D geometric prior and temporal smoothness.
- 531, TITLE: Answer-Type Prediction for Visual Question Answering
http://openaccess.thecvf.com/content_cvpr_2016/html/Kafle_Answer-Type_Prediction_for_CVPR_2016_paper.html
AUTHORS: Kushal Kafle, Christopher Kanan
HIGHLIGHT: In this paper, we build a system capable of answering open-ended text-based questions about images, which is known as Visual Question Answering (VQA).
- 532, TITLE: Visual Word2Vec (vis-w2v): Learning Visually Grounded Word Embeddings Using Abstract Scenes
http://openaccess.thecvf.com/content_cvpr_2016/html/Kottur_Visual_Word2Vec_vis-w2v_CVPR_2016_paper.html
AUTHORS: Satwik Kottur, Ramakrishna Vedantam, Jose M. F. Moura, Devi Parikh
HIGHLIGHT: We propose a model to learn visually grounded word embeddings (vis-w2v) to capture visual notions of semantic relatedness.
- 533, TITLE: Visual7W: Grounded Question Answering in Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhu_Visual7W_Grounded_Question_CVPR_2016_paper.html
AUTHORS: Yuke Zhu, Oliver Groth, Michael Bernstein, Li Fei-Fei
HIGHLIGHT: We establish a semantic link between textual descriptions and image regions by object-level grounding.
- 534, TITLE: Learning Deep Structure-Preserving Image-Text Embeddings
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Learning_Deep_Structure-Preserving_CVPR_2016_paper.html
AUTHORS: Liwei Wang, Yin Li, Svetlana Lazebnik
HIGHLIGHT: This paper proposes a method for learning joint embeddings of images and text using a two-branch neural network with multiple layers of linear projections followed by nonlinearities.
- 535, TITLE: Yin and Yang: Balancing and Answering Binary Visual Questions
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Yin_and_Yang_CVPR_2016_paper.html
AUTHORS: Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, Devi Parikh
HIGHLIGHT: In this paper, we address binary Visual Question Answering (VQA) on abstract scenes.

- 536, TITLE: GIFT: A Real-Time and Scalable 3D Shape Search Engine
http://openaccess.thecvf.com/content_cvpr_2016/html/Bai_GIFT_A_Real-Time_CVPR_2016_paper.html
AUTHORS: Song Bai, Xiang Bai, Zhichao Zhou, Zhaoxiang Zhang, Longin Jan Latecki
HIGHLIGHT: In this paper, we present a real-time 3D shape search engine based on the projective images of 3D shapes.
- 537, TITLE: Functional Faces: Groupwise Dense Correspondence Using Functional Maps
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Functional_Faces_Groupwise_CVPR_2016_paper.html
AUTHORS: Chao Zhang, William A. P. Smith, Arnaud Dessein, Nick Pears, Hang Dai
HIGHLIGHT: In this paper we present a method for computing dense correspondence between a set of 3D face meshes using functional maps.
- 538, TITLE: Similarity Metric For Curved Shapes In Euclidean Space
http://openaccess.thecvf.com/content_cvpr_2016/html/Demisse_Similarity_Metric_For_CVPR_2016_paper.html
AUTHORS: Girum G. Demisse, Djamilia Aouada, Bjorn Ottersten
HIGHLIGHT: In this paper, we introduce a similarity metric for curved shapes that can be described, distinctively, by ordered points.
- 539, TITLE: Shape Analysis With Hyperbolic Wasserstein Distance
http://openaccess.thecvf.com/content_cvpr_2016/html/Shi_Shape_Analysis_With_CVPR_2016_paper.html
AUTHORS: Jie Shi, Wen Zhang, Yalin Wang
HIGHLIGHT: This paper proposes a novel framework to compute Wasserstein distance between general topological surfaces with hyperbolic metric.
- 540, TITLE: Tensor Power Iteration for Multi-Graph Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Shi_Tensor_Power_Iteration_CVPR_2016_paper.html
AUTHORS: Xinchu Shi, Haibin Ling, Weiming Hu, Junliang Xing, Yanning Zhang
HIGHLIGHT: In this work, we address this challenging problem in a principled way under the rank-1 tensor approximation framework.
- 541, TITLE: Multivariate Regression on the Grassmannian for Predicting Novel Domains
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Multivariate_Regression_on_CVPR_2016_paper.html
AUTHORS: Yongxin Yang, Timothy M. Hospedales
HIGHLIGHT: In this paper, we consider the case where domains are parametrised by a vector of continuous values (e.g., time, lighting or view angle).
- 542, TITLE: Learning Cross-Domain Landmarks for Heterogeneous Domain Adaptation
http://openaccess.thecvf.com/content_cvpr_2016/html/Tsai_Learning_Cross-Domain_Landmarks_CVPR_2016_paper.html
AUTHORS: Yao-Hung Hubert Tsai, Yi-Ren Yeh, Yu-Chiang Frank Wang
HIGHLIGHT: In this paper, we propose a novel learning algorithm of Cross-Domain Landmark Selection (CDLS) for solving the above task.
- 543, TITLE: Geospatial Correspondences for Multimodal Registration
http://openaccess.thecvf.com/content_cvpr_2016/html/Marcos_Geospatial_Correspondences_for_CVPR_2016_paper.html
AUTHORS: Diego Marcos, Raffay Hamid, Devis Tuia
HIGHLIGHT: To tackle these challenges, we propose a mid-level sensor-invariant representation that encodes image regions in terms of the spatial distribution of their spectral neighbors.
- 544, TITLE: Constrained Deep Transfer Feature Learning and Its Applications
http://openaccess.thecvf.com/content_cvpr_2016/html/Wu_Constrained_Deep_Transfer_CVPR_2016_paper.html
AUTHORS: Yue Wu, Qiang Ji
HIGHLIGHT: To address these issues, we introduce a constrained deep transfer feature learning method to perform simultaneous transfer learning and feature learning by performing transfer learning in a progressively improving feature space iteratively in order to better narrow the gap between the target domain and the source domain for effective transfer of the data from source domain to target domain.
- 545, TITLE: Deep Canonical Time Warping
http://openaccess.thecvf.com/content_cvpr_2016/html/Trigeorgis_Deep_Canonical_Time_CVPR_2016_paper.html
AUTHORS: George Trigeorgis, Mihalis A. Nicolaou, Stefanos Zafeiriou, Bjorn W. Schuller

HIGHLIGHT: To this end, we present the Deep Canonical Time Warping (DCTW), a method which automatically learns complex non-linear representations of multiple time-series, generated such that (i) they are highly correlated, and (ii) temporally in alignment.

546, **TITLE:** Multilinear Hyperplane Hashing
http://openaccess.thecvf.com/content_cvpr_2016/html/Liu_Multilinear_Hyperplane_Hashing_CVPR_2016_paper.html
AUTHORS: Xianglong Liu, Xinjie Fan, Cheng Deng, Zhujin Li, Hao Su, Dacheng Tao
HIGHLIGHT: To overcome this problem, this paper proposes a multilinear hyperplane hashing that generates a hash bit using multiple linear projections.

547, **TITLE:** Large Scale Hard Sample Mining With Monte Carlo Tree Search
http://openaccess.thecvf.com/content_cvpr_2016/html/Canevet_Large_Scale_Hard_CVPR_2016_paper.html
AUTHORS: Olivier Canevet, Francois Fleuret
HIGHLIGHT: We investigate an efficient strategy to collect false positives from very large training sets in the context of object detection.

548, **TITLE:** Multi-Label Ranking From Positive and Unlabeled Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Kanehira_Multi-Label_Ranking_From_CVPR_2016_paper.html
AUTHORS: Atsushi Kanehira, Tatsuya Harada
HIGHLIGHT: In this paper, we specifically examine the training of a multi-label classifier from data with incompletely assigned labels.

549, **TITLE:** Joint Unsupervised Learning of Deep Representations and Image Clusters
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Joint_Unsupervised_Learning_CVPR_2016_paper.html
AUTHORS: Jianwei Yang, Devi Parikh, Dhruv Batra
HIGHLIGHT: In this paper, we propose a recurrent framework for joint unsupervised learning of deep representations and image clusters.

550, **TITLE:** Kernel Sparse Subspace Clustering on Symmetric Positive Definite Manifolds
http://openaccess.thecvf.com/content_cvpr_2016/html/Yin_Kernel_Sparse_Subspace_CVPR_2016_paper.html
AUTHORS: Ming Yin, Yi Guo, Junbin Gao, Zhaoshui He, Shengli Xie
HIGHLIGHT: In this paper, by embedding the SPD matrices into a Reproducing Kernel Hilbert Space (RKHS), a kernel subspace clustering method is constructed on the SPD manifold through an appropriate Log-Euclidean kernel, termed as kernel sparse subspace clustering on the SPD Riemannian manifold (KSSCR).

551, **TITLE:** Symmetry reCAPTCHA
http://openaccess.thecvf.com/content_cvpr_2016/html/Funk_Symmetry_reCAPTCHA_CVPR_2016_paper.html
AUTHORS: Chris Funk, Yanxi Liu
HIGHLIGHT: Our systematic study reveals significant difference between human labeled (reflection and rotation) symmetries on photos and the output of computer vision algorithms on the same photo set.

552, **TITLE:** Unsupervised Learning of Discriminative Attributes and Visual Representations
http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Unsupervised_Learning_of_CVPR_2016_paper.html
AUTHORS: Chen Huang, Chen Change Loy, Xiaoou Tang
HIGHLIGHT: Given a large unlabeled image collection as input, we train deep Convolutional Neural Networks (CNNs) to output a set of discriminative, binary attributes often with semantic meanings.

553, **TITLE:** When VLAD Met Hilbert
http://openaccess.thecvf.com/content_cvpr_2016/html/Harandi_When_VLAD_Met_CVPR_2016_paper.html
AUTHORS: Mehrtash Harandi, Mathieu Salzmann, Fatih Porikli
HIGHLIGHT: In this paper, we address two fundamental limitations of VLAD: its requirement for the local descriptors to have vector form and its restriction to linear classifiers due to its high-dimensionality.

554, **TITLE:** Approximate Log-Hilbert-Schmidt Distances Between Covariance Operators for Image Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Minh_Approximate_Log-Hilbert-Schmidt_Distances_CVPR_2016_paper.html
AUTHORS: Ha Quang Minh, Marco San Biagio, Loris Bazzani, Vittorio Murino
HIGHLIGHT: This paper presents a novel framework for visual object recognition using infinite-dimensional covariance operators of input features, in the paradigm of kernel methods on infinite-dimensional Riemannian manifolds.

555, **TITLE:** Subspace Clustering With Priors via Sparse Quadratically Constrained Quadratic Programming

- http://openaccess.thecvf.com/content_cvpr_2016/html/Cheng_Subspace_Clustering_With_CVPR_2016_paper.html
AUTHORS: Yongfang Cheng, Yin Wang, Mario Sznaier, Octavia Camps
HIGHLIGHT: This paper considers the problem of recovering a subspace arrangement from noisy samples, potentially corrupted with outliers.
- 556, TITLE: Robust Tensor Factorization With Unknown Noise
http://openaccess.thecvf.com/content_cvpr_2016/html/Chen_Robust_Tensor_Factorization_CVPR_2016_paper.html
AUTHORS: Xi'ai Chen, Zhi Han, Yao Wang, Qian Zhao, Deyu Meng, Yandong Tang
HIGHLIGHT: In this paper, we propose a tensor factorization algorithm to model the noise as a Mixture of Gaussians (MoG).
- 557, TITLE: Kernel Approximation via Empirical Orthogonal Decomposition for Unsupervised Feature Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Mukuta_Kernel_Approximation_via_CVPR_2016_paper.html
AUTHORS: Yusuke Mukuta, Tatsuya Harada
HIGHLIGHT: In this paper, we propose a method that has good generalization performance without high-complexity postprocessing via empirical orthogonal decomposition using the probability distribution estimated from training data.
- 558, TITLE: Active Learning for Delineation of Curvilinear Structures
http://openaccess.thecvf.com/content_cvpr_2016/html/Mosinska-Domanska_Active_Learning_for_CVPR_2016_paper.html
AUTHORS: Agata Mosinska-Domanska, Raphael Sznitman, Przemyslaw Glowacki, Pascal Fua
HIGHLIGHT: In this paper, we propose an Active Learning approach that considerably speeds up the annotation process.
- 559, TITLE: Recognizing Emotions From Abstract Paintings Using Non-Linear Matrix Completion
http://openaccess.thecvf.com/content_cvpr_2016/html/Alameda-Pineda_Recognizing_Emotions_From_CVPR_2016_paper.html
AUTHORS: Xavier Alameda-Pineda, Elisa Ricci, Yan Yan, Nicu Sebe
HIGHLIGHT: In this study we introduce non-linear matrix completion (NLMC), thus extending classical linear matrix completion techniques to the non-linear case.
- 560, TITLE: Tensor Robust Principal Component Analysis: Exact Recovery of Corrupted Low-Rank Tensors via Convex Optimization
http://openaccess.thecvf.com/content_cvpr_2016/html/Lu_Tensor_Robust_Principal_CVPR_2016_paper.html
AUTHORS: Canyi Lu, Jiashi Feng, Yudong Chen, Wei Liu, Zhouchen Lin, Shuicheng Yan
HIGHLIGHT: In this work, we prove that under certain suitable assumptions, we can recover both the low-rank and the sparse components exactly by simply solving a convex program whose objective is a weighted combination of the tensor nuclear norm and the ℓ_1 -norm, i.e., $\min L, E \text{ s.t. } \|L\|_* + \lambda \|E\|_1 \text{ s.t. } X = L + E$.
- 561, TITLE: Sliced Wasserstein Kernels for Probability Distributions
http://openaccess.thecvf.com/content_cvpr_2016/html/Kolouri_Sliced_Wasserstein_Kernels_CVPR_2016_paper.html
AUTHORS: Soheil Kolouri, Yang Zou, Gustavo K. Rohde
HIGHLIGHT: In this paper, we exploit the widely used kernel methods and provide a family of provably positive definite kernels based on the Sliced Wasserstein distance and demonstrate the benefits of these kernels in a variety of learning tasks.
- 562, TITLE: Trace Quotient Meets Sparsity: A Method for Learning Low Dimensional Image Representations
http://openaccess.thecvf.com/content_cvpr_2016/html/Wei_Trace_Quotient_Meets_CVPR_2016_paper.html
AUTHORS: Xian Wei, Hao Shen, Martin Kleinsteuber
HIGHLIGHT: This paper presents an algorithm that allows to learn low dimensional representations of images in an unsupervised manner.
- 563, TITLE: Backtracking ScSPM Image Classifier for Weakly Supervised Top-Down Saliency
http://openaccess.thecvf.com/content_cvpr_2016/html/Cholakkal_Backtracking_ScSPM_Image_CVPR_2016_paper.html
AUTHORS: Hisham Cholakkal, Jubin Johnson, Deepu Rajan
HIGHLIGHT: We propose a weakly supervised top-down saliency framework using only binary labels that indicate the presence/absence of an object in an image.
- 564, TITLE: MSR-VTT: A Large Video Description Dataset for Bridging Video and Language
http://openaccess.thecvf.com/content_cvpr_2016/html/Xu_MSR-VTT_A_Large_CVPR_2016_paper.html
AUTHORS: Jun Xu, Tao Mei, Ting Yao, Yong Rui
HIGHLIGHT: In this paper we present MSR-VTT (standing for "ABC-Video to Text") which is a new large-scale video benchmark for video understanding, especially the emerging task of translating video to text.
- 565, TITLE: NetVLAD: CNN Architecture for Weakly Supervised Place Recognition

http://openaccess.thecvf.com/content_cvpr_2016/html/Arandjelovic_NetVLAD_CNN_Architecture_CVPR_2016_paper.html

AUTHORS: Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, Josef Sivic

HIGHLIGHT: We present the following three principal contributions.

566, TITLE: Structural-RNN: Deep Learning on Spatio-Temporal Graphs

http://openaccess.thecvf.com/content_cvpr_2016/html/Jain_Structural-RNN_Deep_Learning_CVPR_2016_paper.html

AUTHORS: Ashesh Jain, Amir R. Zamir, Silvio Savarese, Ashutosh Saxena

HIGHLIGHT: In this paper, we propose an approach for combining the power of high-level spatio-temporal graphs and sequence learning success of Recurrent Neural Networks (RNNs).

567, TITLE: Learning to Select Pre-Trained Deep Representations With Bayesian Evidence Framework

http://openaccess.thecvf.com/content_cvpr_2016/html/Kim_Learning_to_Select_CVPR_2016_paper.html

AUTHORS: Yong-Deok Kim, Taewoong Jang, Bohyung Han, Seungjin Choi

HIGHLIGHT: We propose a Bayesian evidence framework to facilitate transfer learning from pre-trained deep convolutional neural networks (CNNs).

568, TITLE: Synthesized Classifiers for Zero-Shot Learning

http://openaccess.thecvf.com/content_cvpr_2016/html/Changpinyo_Synthesized_Classifiers_for_CVPR_2016_paper.html

AUTHORS: Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, Fei Sha

HIGHLIGHT: We propose to tackle this problem from the perspective of manifold learning.

To this end, we introduce a set of "phantom" object classes whose coordinates live in both the semantic space and the model space.

569, TITLE: Semi-Supervised Vocabulary-Informed Learning

http://openaccess.thecvf.com/content_cvpr_2016/html/Fu_Semi-Supervised_Vocabulary-Informed_Learning_CVPR_2016_paper.html

AUTHORS: Yanwei Fu, Leonid Sigal

HIGHLIGHT: Specifically, we propose a maximum margin framework for semantic manifold-based recognition that incorporates distance constraints from (both supervised and unsupervised) vocabulary atoms, ensuring that labeled samples are projected closest to their correct prototypes, in the embedding space, than to others.

570, TITLE: Simultaneous Clustering and Model Selection for Tensor Affinities

http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Simultaneous_Clustering_and_CVPR_2016_paper.html

AUTHORS: Zhuwen Li, Shuoguang Yang, Loong-Fah Cheong, Kim-Chuan Toh

HIGHLIGHT: We consider this problem in the domain where the affinity relations involve groups of more than two nodes.

571, TITLE: Discriminatively Embedded K-Means for Multi-View Clustering

http://openaccess.thecvf.com/content_cvpr_2016/html/Xu_Discriminatively_Embedded_K-Means_CVPR_2016_paper.html

AUTHORS: Jinglin Xu, Junwei Han, Feiping Nie

HIGHLIGHT: To address this problem, this paper proposes a novel multi-view clustering method called Discriminatively Embedded K-Means (DEKM), which embeds the synchronous learning of multiple discriminative subspaces into multi-view K-Means clustering to construct a unified framework, and adaptively control the intercoordinations between these subspaces simultaneously.

572, TITLE: Min Norm Point Algorithm for Higher Order MRF-MAP Inference

http://openaccess.thecvf.com/content_cvpr_2016/html/Shanu_Min_Norm_Point_CVPR_2016_paper.html

AUTHORS: Ishant Shanu, Chetan Arora, Parag Singla

HIGHLIGHT: We adapt a well known Min Norm Point algorithm from mathematical optimization literature to exploit the sum of submodular structure found in the MRF-MAP formulation.

573, TITLE: Learning Deep Representation for Imbalanced Classification

http://openaccess.thecvf.com/content_cvpr_2016/html/Huang_Learning_Deep_Representation_CVPR_2016_paper.html

AUTHORS: Chen Huang, Yining Li, Chen Change Loy, Xiaoou Tang

HIGHLIGHT: In this paper, we conduct extensive and systematic experiments to validate the effectiveness of these classic schemes for representation learning on class-imbalanced data.

574, TITLE: Learning Local Image Descriptors With Deep Siamese and Triplet Convolutional Networks by Minimising Global Loss Functions

http://openaccess.thecvf.com/content_cvpr_2016/html/G_Learning_Local_Image_CVPR_2016_paper.html

AUTHORS: Vijay Kumar B G, Gustavo Carneiro, Ian Reid

HIGHLIGHT: Using the UBC benchmark dataset for comparing local image descriptors, we show that the triplet network produces a more accurate embedding than the siamese network in terms of the UBC dataset errors.

- 575, TITLE: Sparse Coding for Third-Order Super-Symmetric Tensor Descriptors With Application to Texture Recognition
http://openaccess.thecvf.com/content_cvpr_2016/html/Koniusz_Sparse_Coding_for_CVPR_2016_paper.html
AUTHORS: Piotr Koniusz, Anoop Chierian
HIGHLIGHT: In this paper, we study third-order super-symmetric tensor descriptors in the context of dictionary learning and sparse coding.
- 576, TITLE: Random Features for Sparse Signal Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Chang_Random_Features_for_CVPR_2016_paper.html
AUTHORS: Jen-Hao Rick Chang, Aswin C. Sankaranarayanan, B. V. K. Vijaya Kumar
HIGHLIGHT: In this paper, we derive performance guarantees for random features on signals, like images, that enjoy sparse representations and show that the number of random features required to achieve a desired approximation of the kernel similarity matrix can be significantly smaller for sparse signals.
- 577, TITLE: High-Quality Depth From Uncalibrated Small Motion Clip
http://openaccess.thecvf.com/content_cvpr_2016/html/Ha_High-Quality_Depth_From_CVPR_2016_paper.html
AUTHORS: Hyowon Ha, Sunghoon Im, Jaesik Park, Hae-Gon Jeon, In So Kweon
HIGHLIGHT: We propose a novel approach that generates a high-quality depth map from a set of images captured with a small viewpoint variation, namely small motion clip.
- 578, TITLE: Efficient 3D Room Shape Recovery From a Single Panorama
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_Efficient_3D_Room_CVPR_2016_paper.html
AUTHORS: Hao Yang, Hui Zhang
HIGHLIGHT: We propose a method to recover the shape of a 3D room from a full-view indoor panorama.
- 579, TITLE: Structured Prediction of Unobserved Voxels From a Single Depth Image
http://openaccess.thecvf.com/content_cvpr_2016/html/Firman_Structured_Prediction_of_CVPR_2016_paper.html
AUTHORS: Michael Firman, Oisín Mac Aodha, Simon Julier, Gabriel J. Brostow
HIGHLIGHT: Exploring this hypothesis, we propose an algorithm that can complete the unobserved geometry of tabletop-sized objects, based on a supervised model trained on already available volumetric elements.
- 580, TITLE: HyperDepth: Learning Depth From Structured Light Without Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Fanello_HyperDepth_Learning_Depth_CVPR_2016_paper.html
AUTHORS: Sean Ryan Fanello, Christoph Rhemann, Vladimir Tankovich, Adarsh Kowdle, Sergio Orts Escolano, David Kim, Shahram Izadi
HIGHLIGHT: We contribute an algorithm for solving this correspondence problem efficiently, without compromising depth accuracy.
- 581, TITLE: SVBRDF-Invariant Shape and Reflectance Estimation From Light-Field Cameras
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_SVBRDF-Invariant_Shape_and_CVPR_2016_paper.html
AUTHORS: Ting-Chun Wang, Manmohan Chandraker, Alexei A. Efros, Ravi Ramamoorthi
HIGHLIGHT: In this paper, we derive a spatially-varying (SV)BRDF-invariant theory for recovering 3D shape and reflectance from light-field cameras.
We present extensive synthetic data on the entire MERL BRDF dataset, as well as a number of real examples to validate the theory, where we simultaneously recover shape and BRDFs from a single image taken with a Lytro Illum camera.
- 582, TITLE: Semantic 3D Reconstruction With Continuous Regularization and Ray Potentials Using a Visibility Consistency Constraint
http://openaccess.thecvf.com/content_cvpr_2016/html/Savinov_Semantic_3D_Reconstruction_CVPR_2016_paper.html
AUTHORS: Nikolay Savinov, Christian Hane, Lubor Ladicky, Marc Pollefeys
HIGHLIGHT: We propose an approach for dense semantic 3D reconstruction which uses a data term that is defined as potentials over viewing rays, combined with continuous surface area penalization.
- 583, TITLE: Theory and Practice of Structure-From-Motion Using Affine Correspondences
http://openaccess.thecvf.com/content_cvpr_2016/html/Raposo_Theory_and_Practice_CVPR_2016_paper.html
AUTHORS: Carolina Raposo, Joao P. Barreto
HIGHLIGHT: This article is a step forward into this direction, by providing a clear account about how ACs constrain the two-view geometry, and by proposing new algorithms for plane segmentation and visual odometry that compare favourably with respect to methods relying in PCs.

- 584, TITLE: Just Look at the Image: Viewpoint-Specific Surface Normal Prediction for Improved Multi-View Reconstruction
http://openaccess.thecvf.com/content_cvpr_2016/html/Galliani_Just_Look_at_CVPR_2016_paper.html
AUTHORS: Silvano Galliani, Konrad Schindler
HIGHLIGHT: We present a multi-view reconstruction method that combines conventional multi-view stereo (MVS) with appearance-based normal prediction, to obtain dense and accurate 3D surface models.
- 585, TITLE: From Dusk Till Dawn: Modeling in the Dark
http://openaccess.thecvf.com/content_cvpr_2016/html/Radenovic_From_Dusk_Till_CVPR_2016_paper.html
AUTHORS: Filip Radenovic, Johannes L. Schonberger, Dinghuang Ji, Jan-Michael Frahm, Ondrej Chum, Jiri Matas
HIGHLIGHT: We present an algorithm that leverages the appearance variety to obtain more complete and accurate scene geometry along with consistent multi-illumination appearance information.
- 586, TITLE: Accelerated Generative Models for 3D Point Cloud Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Eckart_Accelerated_Generative_Models_CVPR_2016_paper.html
AUTHORS: Benjamin Eckart, Kihwan Kim, Alejandro Troccoli, Alonzo Kelly, Jan Kautz
HIGHLIGHT: In this paper we introduce a method for constructing compact generative representations of PCD at multiple levels of detail.
- 587, TITLE: Monocular Depth Estimation Using Neural Regression Forest
http://openaccess.thecvf.com/content_cvpr_2016/html/Roy_Monocular_Depth_Estimation_CVPR_2016_paper.html
AUTHORS: Anirban Roy, Sinisa Todorovic
HIGHLIGHT: This paper presents a novel deep architecture, called neural regression forest (NRF), for depth estimation from a single image.
- 588, TITLE: DeepStereo: Learning to Predict New Views From the World's Imagery
http://openaccess.thecvf.com/content_cvpr_2016/html/Flynn_DeepStereo_Learning_to_CVPR_2016_paper.html
AUTHORS: John Flynn, Ivan Neulander, James Philbin, Noah Snavely
HIGHLIGHT: In this work, we present a novel deep architecture that performs new view synthesis directly from pixels, trained from a large number of posed image sets.
- 589, TITLE: WIDER FACE: A Face Detection Benchmark
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_WIDER_FACE_A_CVPR_2016_paper.html
AUTHORS: Shuo Yang, Ping Luo, Chen-Change Loy, Xiaoou Tang
HIGHLIGHT: To facilitate future face detection research, we introduce the WIDER FACE dataset, which is 10 times larger than existing datasets.
- 590, TITLE: Situation Recognition: Visual Semantic Role Labeling for Image Understanding
http://openaccess.thecvf.com/content_cvpr_2016/html/Yatskar_Situation_Recognition_Visual_CVPR_2016_paper.html
AUTHORS: Mark Yatskar, Luke Zettlemoyer, Ali Farhadi
HIGHLIGHT: This paper introduces situation recognition, the problem of producing a concise summary of the situation an image depicts including: (1) the main activity (e.g., clipping), (2) the participating actors, objects, substances, and locations (e.g., man, shears, sheep, wool, and field) and most importantly (3) the roles these participants play in the activity (e.g., the man is clipping, the shears are his tool, the wool is being clipped from the sheep, and the clipping is in a field).
- 591, TITLE: A 3D Morphable Model Learnt From 10,000 Faces
http://openaccess.thecvf.com/content_cvpr_2016/html/Booth_A_3D_Morphable_CVPR_2016_paper.html
AUTHORS: James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, David Dunaway
HIGHLIGHT: We present Large Scale Facial Model (LSFM) -- a 3D Morphable Model (3DMM) automatically constructed from 9,663 distinct facial identities.
- 592, TITLE: Some Like It Hot - Visual Guidance for Preference Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Rothe_Some_Like_It_CVPR_2016_paper.html
AUTHORS: Rasmus Rothe, Radu Timofte, Luc Van Gool
HIGHLIGHT: We demonstrate our algorithms on howhot.io which went viral around the Internet with more than 50 million pictures evaluated in the first month.
- 593, TITLE: EmotioNet: An Accurate, Real-Time Algorithm for the Automatic Annotation of a Million Facial Expressions in the Wild
http://openaccess.thecvf.com/content_cvpr_2016/html/Benitez-Quiroz_EmotioNet_An_Accurate_CVPR_2016_paper.html
AUTHORS: C. Fabian Benitez-Quiroz, Ramprakash Srinivasan, Aleix M. Martinez

HIGHLIGHT: Herein, we present a novel computer vision algorithm to annotate a large database of one million images of facial expressions of emotion in the wild (i.e., face images downloaded from the Internet).

594, **TITLE:** ForgetMeNot: Memory-Aware Forensic Facial Sketch Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Ouyang_ForgetMeNot_Memory-Aware_Forensic_CVPR_2016_paper.html
AUTHORS: Shuxin Ouyang, Timothy M. Hospedales, Yi-Zhe Song, Xueming Li
HIGHLIGHT: In this paper we address the memory problem head on by introducing a database of 400 forensic sketches created at different time-delays.

595, **TITLE:** LOMo: Latent Ordinal Model for Facial Analysis in Videos
http://openaccess.thecvf.com/content_cvpr_2016/html/Sikka_LOMo_Latent_Ordinal_CVPR_2016_paper.html
AUTHORS: Karan Sikka, Gaurav Sharma, Marian Bartlett
HIGHLIGHT: We study the problem of facial analysis in videos.

596, **TITLE:** Discriminative Invariant Kernel Features: A Bells-and-Whistles-Free Approach to Unsupervised Face Recognition and Pose Estimation
http://openaccess.thecvf.com/content_cvpr_2016/html/Pal_Discriminative_Invariant_Kernel_CVPR_2016_paper.html
AUTHORS: Dipan K. Pal, Felix Juefei-Xu, Marios Savvides
HIGHLIGHT: We propose an explicitly discriminative and 'simple' approach to generate invariance to nuisance transformations modeled as unitary.

597, **TITLE:** Bottom-Up and Top-Down Reasoning With Hierarchical Rectified Gaussians
http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Bottom-Up_and_Top-Down_CVPR_2016_paper.html
AUTHORS: Peiyun Hu, Deva Ramanan
HIGHLIGHT: Such approaches tend to work in a "unidirectional" bottom-up feed-forward fashion.

598, **TITLE:** Fits Like a Glove: Rapid and Reliable Hand Shape Personalization
http://openaccess.thecvf.com/content_cvpr_2016/html/Tan_Fits_Like_a_CVPR_2016_paper.html
AUTHORS: David Joseph Tan, Thomas Cashman, Jonathan Taylor, Andrew Fitzgibbon, Daniel Tarlow, Sameh Khamis, Shahram Izadi, Jamie Shotton
HIGHLIGHT: We present a fast, practical method for personalizing a hand shape basis to an individual user's detailed hand shape using only a small set of depth images.

599, **TITLE:** Slicing Convolutional Neural Network for Crowd Video Understanding
http://openaccess.thecvf.com/content_cvpr_2016/html/Shao_Slicing_Convolutional_Neural_CVPR_2016_paper.html
AUTHORS: Jing Shao, Chen-Change Loy, Kai Kang, Xiaogang Wang
HIGHLIGHT: In this study, we propose a novel spatio-temporal CNN, named Slicing CNN (S-CNN), based on the decomposition of 3D feature maps into 2D spatio- and 2D temporal-slices representations.

600, **TITLE:** Linear Shape Deformation Models With Local Support Using Graph-Based Structured Matrix Factorisation
http://openaccess.thecvf.com/content_cvpr_2016/html/Bernard_Linear_Shape_Deformation_CVPR_2016_paper.html
AUTHORS: Florian Bernard, Peter Gemmar, Frank Hertel, Jorge Goncalves, Johan Thunberg
HIGHLIGHT: Based on matrix factorisation with sparsity and graph-based regularisation terms, we present a method to obtain deformation factors with local support.

601, **TITLE:** Motion From Structure (MfS): Searching for 3D Objects in Cluttered Point Trajectories
http://openaccess.thecvf.com/content_cvpr_2016/html/Vongkulbhisal_Motion_From_Structure_CVPR_2016_paper.html
AUTHORS: Jayakorn Vongkulbhisal, Ricardo Cabral, Fernando De la Torre, Joao P. Costeira
HIGHLIGHT: In this work, we explore the use of 3D shape models to detect objects in videos in an unsupervised manner.

602, **TITLE:** Volumetric and Multi-View CNNs for Object Classification on 3D Data
http://openaccess.thecvf.com/content_cvpr_2016/html/Qi_Volumetric_and_Multi-View_CVPR_2016_paper.html
AUTHORS: Charles R. Qi, Hao Su, Matthias Niessner, Angela Dai, Mengyuan Yan, Leonidas J. Guibas
HIGHLIGHT: In this paper, we aim to improve both volumetric CNNs and multi-view CNNs according to extensive analysis of existing approaches.

603, **TITLE:** Detecting Vanishing Points Using Global Image Context in a Non-Manhattan World
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhai_Detecting_Vanishing_Points_CVPR_2016_paper.html
AUTHORS: Menghua Zhai, Scott Workman, Nathan Jacobs

HIGHLIGHT: We propose a novel method for detecting horizontal vanishing points and the zenith vanishing point in man-made environments.
Our method reverses this process: we propose a set of horizon line candidates and score each based on the vanishing points it contains.

604, **TITLE:** Learning Weight Uncertainty With Stochastic Gradient MCMC for Shape Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Li_Learning_Weight_Uncertainty_CVPR_2016_paper.html
AUTHORS: Chunyuan Li, Andrew Stevens, Changyou Chen, Yunchen Pu, Zhe Gan, Lawrence Carin
HIGHLIGHT: This paper leverages recent advances in stochastic gradient Markov Chain Monte Carlo (SG-MCMC) to learn weight uncertainty in DNNs.

605, **TITLE:** A Field Model for Repairing 3D Shapes
http://openaccess.thecvf.com/content_cvpr_2016/html/Nguyen_A_Field_Model_CVPR_2016_paper.html
AUTHORS: Duc Thanh Nguyen, Binh-Son Hua, Khoi Tran, Quang-Hieu Pham, Sai-Kit Yeung
HIGHLIGHT: This paper proposes a field model for repairing 3D shapes constructed from multi-view RGB data.

606, **TITLE:** GOGMA: Globally-Optimal Gaussian Mixture Alignment
http://openaccess.thecvf.com/content_cvpr_2016/html/Campbell_GOGMA_Globally-Optimal_Gaussian_CVPR_2016_paper.html
AUTHORS: Dylan Campbell, Lars Petersson
HIGHLIGHT: This paper presents the first globally-optimal solution to the 3D rigid Gaussian mixture alignment problem under the L2 distance between mixtures.

607, **TITLE:** Efficient Deep Learning for Stereo Matching
http://openaccess.thecvf.com/content_cvpr_2016/html/Luo_Efficient_Deep_Learning_CVPR_2016_paper.html
AUTHORS: Wenjie Luo, Alexander G. Schwing, Raquel Urtasun
HIGHLIGHT: In contrast, in this paper we propose a matching network which is able to produce very accurate results in less than a second of GPU computation.

608, **TITLE:** Efficient Coarse-To-Fine PatchMatch for Large Displacement Optical Flow
http://openaccess.thecvf.com/content_cvpr_2016/html/Hu_Efficient_Coarse-To-Fine_PatchMatch_CVPR_2016_paper.html
AUTHORS: Yinlin Hu, Rui Song, Yunsong Li
HIGHLIGHT: In this paper we present a simple but powerful matching method works in a coarse-to-fine scheme for optical flow estimation.

609, **TITLE:** FANNG: Fast Approximate Nearest Neighbour Graphs
http://openaccess.thecvf.com/content_cvpr_2016/html/Harwood_FANNG_Fast_Approximate_CVPR_2016_paper.html
AUTHORS: Ben Harwood, Tom Drummond
HIGHLIGHT: We present a new method for approximate nearest neighbour search on large datasets of high dimensional feature vectors, such as SIFT or GIST descriptors.

610, **TITLE:** Exemplar-Driven Top-Down Saliency Detection via Deep Association
http://openaccess.thecvf.com/content_cvpr_2016/html/He_Exemplar-Driven_Top-Down_Saliency_CVPR_2016_paper.html
AUTHORS: Shengfeng He, Rynson W.H. Lau, Qingxiong Yang
HIGHLIGHT: In contrast, we propose in this paper a locate-by-exemplar strategy.

611, **TITLE:** Unconstrained Salient Object Detection via Proposal Subset Optimization
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Unconstrained_Salient_Object_CVPR_2016_paper.html
AUTHORS: Jianming Zhang, Stan Sclaroff, Zhe Lin, Xiaohui Shen, Brian Price, Radomir Mech
HIGHLIGHT: Based on the Maximum a Posteriori principle, we propose a novel subset optimization framework to generate a compact set of detection windows out of noisy proposals.

612, **TITLE:** Recombinator Networks: Learning Coarse-To-Fine Feature Aggregation
http://openaccess.thecvf.com/content_cvpr_2016/html/Honari_Recombinator_Networks_Learning_CVPR_2016_paper.html
AUTHORS: Sina Honari, Jason Yosinski, Pascal Vincent, Christopher Pal
HIGHLIGHT: Here we introduce another model --- dubbed Recombinator Networks --- where coarse features inform finer features early in their formation such that finer features can make use of several layers of computation in deciding how to use coarse features.

613, **TITLE:** End-To-End Saliency Mapping via Probability Distribution Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Jetley_End-To-End_Saliency_Mapping_CVPR_2016_paper.html
AUTHORS: Saumya Jetley, Naila Murray, Eleonora Vig

- HIGHLIGHT:** In this work, we introduce a new saliency map model which formulates a map as a generalized Bernoulli distribution.
- 614, **TITLE:** A Paradigm for Building Generalized Models of Human Image Perception Through Data Fusion
http://openaccess.thecvf.com/content_cvpr_2016/html/Fan_A_Paradigm_for_CVPR_2016_paper.html
AUTHORS: Shaojing Fan, Tian-Tsong Ng, Bryan L. Koenig, Ming Jiang, Qi Zhao
HIGHLIGHT: In this paper, we present a paradigm for building generalized and expandable models of human image perception.
- 615, **TITLE:** Longitudinal Face Modeling via Temporal Deep Restricted Boltzmann Machines
http://openaccess.thecvf.com/content_cvpr_2016/html/Duong_Longitudinal_Face_Modeling_CVPR_2016_paper.html
AUTHORS: Chi Nhan Duong, Khoa Luu, Kha Gia Quach, Tien D. Bui
HIGHLIGHT: This paper presents a deep model approach for face age progression that can efficiently capture the non-linear aging process and automatically synthesize a series of age-progressed faces in various age ranges.
- 616, **TITLE:** Saliency Unified: A Deep Architecture for Simultaneous Eye Fixation Prediction and Salient Object Segmentation
http://openaccess.thecvf.com/content_cvpr_2016/html/Kruthiventi_Saliency_Unified_A_CVPR_2016_paper.html
AUTHORS: Srinivas S. S. Kruthiventi, Vennela Gudisa, Jaley H. Dholakiya, R. Venkatesh Babu
HIGHLIGHT: In this work, we propose a deep convolutional neural network (CNN) capable of predicting eye fixations and segmenting salient objects in a unified framework.
- 617, **TITLE:** Estimating Correspondences of Deformable Objects "In-The-Wild"
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhou_Estimating_Correspondences_of_CVPR_2016_paper.html
AUTHORS: Yuxiang Zhou, Epameinondas Antonakos, Joan Alabort-i-Medina, Anastasios Roussos, Stefanos Zafeiriou
HIGHLIGHT: In this paper, we show for the first time, to the best of our knowledge, that it is possible to construct SDMs by putting object shapes in dense correspondence.
- 618, **TITLE:** Gravitational Approach for Point Set Registration
http://openaccess.thecvf.com/content_cvpr_2016/html/Golyanik_Gravitational_Approach_for_CVPR_2016_paper.html
AUTHORS: Vladislav Golyanik, Sk Aziz Ali, Didier Stricker
HIGHLIGHT: In this paper a new astrodynamics inspired rigid point set registration algorithm is introduced -- the Gravitational Approach (GA).
- 619, **TITLE:** Context-Aware Gaussian Fields for Non-Rigid Point Set Registration
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Context-Aware_Gaussian_Fields_CVPR_2016_paper.html
AUTHORS: Gang Wang, Zhicheng Wang, Yufei Chen, Qiangqiang Zhou, Weidong Zhao
HIGHLIGHT: This paper proposes context-aware Gaussian fields (CA-LapGF) for non-rigid PSR subject to global rigid and local non-rigid geometric constraints, where a laplacian regularized term is added to preserve the intrinsic geometry of the transformed set.
- 620, **TITLE:** Trust No One: Low Rank Matrix Factorization Using Hierarchical RANSAC
http://openaccess.thecvf.com/content_cvpr_2016/html/Oskarsson_Trust_No_One_CVPR_2016_paper.html
AUTHORS: Magnus Oskarsson, Kenneth Batstone, Kalle Astrom
HIGHLIGHT: In this paper we present a system for performing low rank matrix factorization.
- 621, **TITLE:** Relaxation-Based Preprocessing Techniques for Markov Random Field Inference
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Relaxation-Based_Preprocessing_Techniques_CVPR_2016_paper.html
AUTHORS: Chen Wang, Ramin Zabih
HIGHLIGHT: We describe an efficient algorithm to correctly label a subset of the variables for arbitrary MRFs, with particularly good performance on binary MRFs.
- 622, **TITLE:** Sparse Coding for Classification via Discrimination Ensemble
http://openaccess.thecvf.com/content_cvpr_2016/html/Quan_Sparse_Coding_for_CVPR_2016_paper.html
AUTHORS: Yuhui Quan, Yong Xu, Yuping Sun, Yan Huang, Hui Ji
HIGHLIGHT: In this paper, we proposed a discriminative sparse coding method which jointly learns a dictionary for sparse coding and an ensemble classifier for discrimination.
- 623, **TITLE:** Principled Parallel Mean-Field Inference for Discrete Random Fields
http://openaccess.thecvf.com/content_cvpr_2016/html/Baque_Principled_Parallel_Mean-Field_CVPR_2016_paper.html

- AUTHORS: Pierre Baque, Timur Bagautdinov, Francois Fleuret, Pascal Fua
HIGHLIGHT: In this paper, we propose a novel proximal gradient-based approach to optimizing the variational objective.
- 624, TITLE: Guaranteed Outlier Removal With Mixed Integer Linear Programs
http://openaccess.thecvf.com/content_cvpr_2016/html/Chin_Guaranteed_Outlier_Removal_CVPR_2016_paper.html
AUTHORS: Tat-Jun Chin, Yang Heng Kee, Anders Eriksson, Frank Neumann
HIGHLIGHT: Towards the goal of solving maximum consensus exactly, we present guaranteed outlier removal as a technique to reduce the runtime of exact algorithms.
- 625, TITLE: Memory Efficient Max Flow for Multi-Label Submodular MRFs
http://openaccess.thecvf.com/content_cvpr_2016/html/Ajanthan_Memory_Efficient_Max_CVPR_2016_paper.html
AUTHORS: Thalaiyasingam Ajanthan, Richard Hartley, Mathieu Salzmann
HIGHLIGHT: In this paper, we introduce a variant of the max-flow algorithm that requires much less storage.
- 626, TITLE: Proximal Riemannian Pursuit for Large-Scale Trace-Norm Minimization
http://openaccess.thecvf.com/content_cvpr_2016/html/Tan_Proximal_Riemannian_Pursuit_CVPR_2016_paper.html
AUTHORS: Mingkui Tan, Shijie Xiao, Junbin Gao, Dong Xu, Anton van den Hengel, Qinfeng Shi
HIGHLIGHT: In this paper, we propose a proximal Riemannian pursuit (PRP) paradigm which addresses a sequence of trace-norm regularized subproblems defined on nonlinear matrix varieties.
- 627, TITLE: Minimizing the Maximal Rank
http://openaccess.thecvf.com/content_cvpr_2016/html/Bylow_Minimizing_the_Maximal_CVPR_2016_paper.html
AUTHORS: Erik Bylow, Carl Olsson, Fredrik Kahl, Mikael Nilsson
HIGHLIGHT: In this paper a new convex envelope is derived which takes all sub-matrices into account simultaneously.
- 628, TITLE: Solving Temporal Puzzles
http://openaccess.thecvf.com/content_cvpr_2016/html/Dicle_Solving_Temporal_Puzzles_CVPR_2016_paper.html
AUTHORS: Caglayan Dicle, Burak Yilmaz, Octavia Camps, Mario Szaier
HIGHLIGHT: In this paper, based on this intuition, we propose an algorithm for solving time-sort temporal puzzles, defined as scrambled time series that need to be sorted out.
- 629, TITLE: Estimating Sparse Signals With Smooth Support via Convex Programming and Block Sparsity
http://openaccess.thecvf.com/content_cvpr_2016/html/Shah_Estimating_Sparse_Signals_CVPR_2016_paper.html
AUTHORS: Sohil Shah, Tom Goldstein, Christoph Studer
HIGHLIGHT: In this paper, we explore the use of new block l_1 -norm regularizers, which enforce image sparsity while simultaneously promoting smooth support structure.
- 630, TITLE: TenSR: Multi-Dimensional Tensor Sparse Representation
http://openaccess.thecvf.com/content_cvpr_2016/html/Qi_TenSR_Multi-Dimensional_Tensor_CVPR_2016_paper.html
AUTHORS: Na Qi, Yunhui Shi, Xiaoyan Sun, Baocai Yin
HIGHLIGHT: In this paper, we propose a new sparse model TenSR based on tensor for MD data representation along with the corresponding MD sparse coding and MD dictionary learning algorithms.
- 631, TITLE: Moral Lineage Tracing
http://openaccess.thecvf.com/content_cvpr_2016/html/Jug_Moral_Lineage_Tracing_CVPR_2016_paper.html
AUTHORS: Florian Jug, Evgeny Levinkov, Corinna Blasse, Eugene W. Myers, Bjoern Andres
HIGHLIGHT: We propose an integer linear program (ILP) whose feasible solutions define, for every image in a sequence, a decomposition into cells (segmentation) and, across images, a lineage forest of cells (tracing).
- 632, TITLE: Globally Optimal Rigid Intensity Based Registration: A Fast Fourier Domain Approach
http://openaccess.thecvf.com/content_cvpr_2016/html/Nasihatkon_Globally_Optimal_Rigid_CVPR_2016_paper.html
AUTHORS: Behrooz Nasihatkon, Frida Fejne, Fredrik Kahl
HIGHLIGHT: In this paper, we propose a dual algorithm in which the optimization is done in the Fourier domain, and multiple resolution levels are replaced by multiple frequency bands.
- 633, TITLE: On Benefits of Selection Diversity via Bilevel Exclusive Sparsity
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_On_Benefits_of_CVPR_2016_paper.html
AUTHORS: Haichuan Yang, Yijun Huang, Lam Tran, Ji Liu, Shuai Huang
HIGHLIGHT: In this paper, we proposed a general bilevel exclusive sparsity formulation to pursue the diversity by restricting the overall sparsity and the sparsity in each group.

- 634, TITLE: Fast Training of Triplet-Based Deep Binary Embedding Networks
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhuang_Fast_Training_of_CVPR_2016_paper.html
AUTHORS: Bohan Zhuang, Guosheng Lin, Chunhua Shen, Ian Reid
HIGHLIGHT: In this paper, we aim to learn a mapping (or embedding) from images to a compact binary space in which Hamming distances correspond to a ranking measure for the image retrieval task.
- 635, TITLE: Marr Revisited: 2D-3D Alignment via Surface Normal Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Bansal_Marr_Revisited_2D-3D_CVPR_2016_paper.html
AUTHORS: Aayush Bansal, Bryan Russell, Abhinav Gupta
HIGHLIGHT: We introduce an approach that leverages surface normal predictions, along with appearance cues, to retrieve 3D models for objects depicted in 2D still images from a large CAD object library.
- 636, TITLE: Recovering the Missing Link: Predicting Class-Attribute Associations for Unsupervised Zero-Shot Learning
http://openaccess.thecvf.com/content_cvpr_2016/html/Al-Halah_Recovering_the_Missing_CVPR_2016_paper.html
AUTHORS: Ziad Al-Halah, Makarand Tapaswi, Rainer Stiefelwagen
HIGHLIGHT: In this work, we aim to carry out attribute-based zero-shot classification in an unsupervised manner.
- 637, TITLE: Fast Zero-Shot Image Tagging
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Fast_Zero-Shot_Image_CVPR_2016_paper.html
AUTHORS: Yang Zhang, Boqing Gong, Mubarak Shah
HIGHLIGHT: We study a particular image-word relevance relation in this paper.
- 638, TITLE: Modality and Component Aware Feature Fusion For RGB-D Scene Classification
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_Modality_and_Component_CVPR_2016_paper.html
AUTHORS: Anran Wang, Jianfei Cai, Jiwen Lu, Tat-Jen Cham
HIGHLIGHT: In this paper, we investigate a framework allowing greater spatial flexibility, in which the Fisher vector (FV) encoded distribution of local CNN features, obtained from a multitude of region proposals per image, is considered instead.
- 639, TITLE: PPP: Joint Pointwise and Pairwise Image Label Prediction
http://openaccess.thecvf.com/content_cvpr_2016/html/Wang_PPP_Joint_Pointwise_CVPR_2016_paper.html
AUTHORS: Yilin Wang, Suhang Wang, Jiliang Tang, Huan Liu, Baoxin Li
HIGHLIGHT: In particular, we provide a principled way to capture the relations between class labels, tags and attributes; and propose a novel framework PPP(Pointwise and Pairwise image label Prediction), which is based on overlapped group structure extracted from the pointwise-pairwise-label bipartite graph.
- 640, TITLE: Cataloging Public Objects Using Aerial and Street-Level Images - Urban Trees
http://openaccess.thecvf.com/content_cvpr_2016/html/Wegner_Cataloging_Public_Objects_CVPR_2016_paper.html
AUTHORS: Jan D. Wegner, Steven Branson, David Hall, Konrad Schindler, Pietro Perona
HIGHLIGHT: We introduce a solution that adapts state-of-the-art CNN-based object detectors and classifiers.
- 641, TITLE: Deep Exemplar 2D-3D Detection by Adapting From Real to Rendered Views
http://openaccess.thecvf.com/content_cvpr_2016/html/Massa_Deep_Exemplar_2D-3D_CVPR_2016_paper.html
AUTHORS: Francisco Massa, Bryan C. Russell, Mathieu Aubry
HIGHLIGHT: This paper presents an end-to-end convolutional neural network (CNN) for 2D-3D exemplar detection.
- 642, TITLE: Zero-Shot Learning via Joint Latent Similarity Embedding
http://openaccess.thecvf.com/content_cvpr_2016/html/Zhang_Zero-Shot_Learning_via_CVPR_2016_paper.html
AUTHORS: Ziming Zhang, Venkatesh Saligrama
HIGHLIGHT: We formulate ZSR as a binary prediction problem.
- 643, TITLE: CRAFT Objects From Images
http://openaccess.thecvf.com/content_cvpr_2016/html/Yang_CRAFT_Objects_From_CVPR_2016_paper.html
AUTHORS: Bin Yang, Junjie Yan, Zhen Lei, Stan Z. Li
HIGHLIGHT: In this paper, we push the "divide and conquer" solution even further by dividing each task into two sub-tasks.