

1, TITLE: BA-Net: Dense Bundle Adjustment Networks  
<https://openreview.net/forum?id=B1gabRcYX>  
AUTHORS: Chengzhou Tang, Ping Tan  
TL;DR: This paper introduces a network architecture to solve the structure-from-motion (SfM) problem via feature bundle adjustment (BA)  
HIGHLIGHT: This paper introduces a network architecture to solve the structure-from-motion (SfM) problem via feature-metric bundle adjustment (BA), which explicitly enforces multi-view geometry constraints in the form of feature-metric error.

2, TITLE: Deterministic Variational Inference for Robust Bayesian Neural Networks  
<https://openreview.net/forum?id=B1l08oAct7>  
AUTHORS: Anqi Wu, Sebastian Nowozin, Edward Meeds, Richard E. Turner, Jos? Miguel Hernandez-Lobato, Alexander L. Gaunt  
TL;DR: A method for eliminating gradient variance and automatically tuning priors for effective training of bayesian neural networks  
HIGHLIGHT: On the application of heteroscedastic regression we demonstrate good predictive performance over alternative approaches.

3, TITLE: Ordered Neurons: Integrating Tree Structures into Recurrent Neural Networks  
<https://openreview.net/forum?id=B1l6qiR5F7>  
AUTHORS: Yikang Shen, Shawn Tan, Alessandro Sordoni, Aaron Courville  
TL;DR: We introduce a new inductive bias that integrates tree structures in recurrent neural networks.  
HIGHLIGHT: This paper proposes to add such inductive bias by ordering the neurons; a vector of master input and forget gates ensures that when a given neuron is updated, all the neurons that follow it in the ordering are also updated.

4, TITLE: Large Scale GAN Training for High Fidelity Natural Image Synthesis  
<https://openreview.net/forum?id=B1xsqj09Fm>  
AUTHORS: Andrew Brock, Jeff Donahue, Karen Simonyan  
TL;DR: GANs benefit from scaling up.  
HIGHLIGHT: To this end, we train Generative Adversarial Networks at the largest scale yet attempted, and study the instabilities specific to such scale.

5, TITLE: Learning deep representations by mutual information estimation and maximization  
<https://openreview.net/forum?id=Bklr3j0cKX>  
AUTHORS: R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, Yoshua Bengio  
TL;DR: We learn deep representation by maximizing mutual information, leveraging structure in the objective, and are able to compute with fully supervised classifiers with comparable architectures  
HIGHLIGHT: This work investigates unsupervised learning of representations by maximizing mutual information between an input and the output of a deep neural network encoder.

6, TITLE: KnockoffGAN: Generating Knockoffs for Feature Selection using Generative Adversarial Networks  
<https://openreview.net/forum?id=ByeZ5jC5YQ>  
AUTHORS: James Jordon, Jinsung Yoon, Mihaela van der Schaar  
TL;DR: n/a  
HIGHLIGHT: In this work, we build on the promising Knockoff framework by developing a flexible knockoff generation model.

7, TITLE: Learning Protein Structure with a Differentiable Simulator  
<https://openreview.net/forum?id=Byg3y3C9Km>  
AUTHORS: John Ingraham, Adam Riesselman, Chris Sander, Debora Marks  
TL;DR: We use an unrolled simulator as an end-to-end differentiable model of protein structure and show it can (sometimes) hierarchically generalize to unseen fold topologies.  
HIGHLIGHT: In this work we aim to bridge the gap between the expressive capacity of energy functions and the practical capabilities of their simulators by using an unrolled Monte Carlo simulation as a model for data.

8, TITLE: ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness  
<https://openreview.net/forum?id=Bygh9j09KX>  
AUTHORS: Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, Wieland Brendel  
TL;DR: ImageNet-trained CNNs are biased towards object texture (instead of shape like humans). Overcoming this major difference between human and machine vision yields improved detection performance and previously unseen robustness to image distortions.  
HIGHLIGHT: We show that ImageNet-trained CNNs are strongly biased towards recognising textures rather than shapes, which is in stark contrast to human behavioural evidence and reveals fundamentally different classification strategies.

9, TITLE: Smoothing the Geometry of Probabilistic Box Embeddings  
<https://openreview.net/forum?id=H1xSNiRcF7>  
AUTHORS: Xiang Li, Luke Vilnis, Dongxu Zhang, Michael Boratko, Andrew McCallum  
TL;DR: Improve hierarchical embedding models using kernel smoothing  
HIGHLIGHT: In this work, we present a novel hierarchical embedding model, inspired by a relaxation of box embeddings into parameterized density functions using Gaussian convolutions over the boxes.

10, TITLE: On Random Deep Weight-Tied Autoencoders: Exact Asymptotic Analysis, Phase Transitions, and Implications to Training  
<https://openreview.net/forum?id=HJx54i05tX>  
AUTHORS: Ping Li, Phan-Minh Nguyen  
TL;DR: We study the behavior of weight-tied multilayer vanilla autoencoders under the assumption of random weights. Via an exact characterization in the limit of large dimensions, our analysis reveals interesting phase transition phenomena.  
HIGHLIGHT: We study the behavior of weight-tied multilayer vanilla autoencoders under the assumption of random weights.

11, TITLE: Meta-Learning Update Rules for Unsupervised Representation Learning  
<https://openreview.net/forum?id=HkNDsiC9KQ>  
AUTHORS: Luke Metz, Niru Maheswaranathan, Brian Cheung, Jascha Sohl-Dickstein  
TL;DR: We learn an unsupervised learning algorithm that produces useful representations from a set of supervised tasks. At test-time, we apply this algorithm to new tasks without any supervision and show performance comparable to a VAE.  
HIGHLIGHT: In this work, we propose instead to directly target later desired tasks by meta-learning an unsupervised learning rule which leads to representations useful for those tasks.

12, TITLE: Transferring Knowledge across Learning Processes  
<https://openreview.net/forum?id=HygBZnRctX>  
AUTHORS: Sebastian Flennerhag, Pablo G. Moreno, Neil D. Lawrence, Andreas Damianou

TL;DR: We propose Leap, a framework that transfers knowledge across learning processes by minimizing the expected distance the training process travels on a task's loss surface.

HIGHLIGHT: We propose Leap, a framework that achieves this by transferring knowledge across learning processes.

13, TITLE: GENERATING HIGH FIDELITY IMAGES WITH SUBSCALE PIXEL NETWORKS AND MULTIDIMENSIONAL UPSCALING

<https://openreview.net/forum?id=HylzTiC5Km>

AUTHORS: Jacob Menick, Nal Kalchbrenner

TL;DR: We show that autoregressive models can generate high fidelity images.

HIGHLIGHT: To address the former challenge, we propose the Subscale Pixel Network (SPN), a conditional decoder architecture that generates an image as a sequence of image slices of equal size.

14, TITLE: Temporal Difference Variational Auto-Encoder

<https://openreview.net/forum?id=S1x4ghC9tQ>

AUTHORS: Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, Theophane Weber

TL;DR: Generative model of temporal data, that builds online belief state, operates in latent space, does jumpy predictions and rollouts of states.

HIGHLIGHT: Motivated by the absence of a model satisfying all these requirements, we propose TD-VAE, a generative sequence model that learns representations containing explicit beliefs about states several steps into the future, and that can be rolled out directly without single-step transitions.

15, TITLE: A Unified Theory of Early Visual Representations from Retina to Cortex through Anatomically Constrained Deep CNNs

<https://openreview.net/forum?id=S1xq3oR5tQ>

AUTHORS: Jack Lindsey, Samuel A. Ocko, Surya Ganguli, Stephane Deny

TL;DR: We reproduced neural representations found in biological visual systems by simulating their neural resource constraints in a deep convolutional model.

HIGHLIGHT: Here, using a deep convolutional neural network trained on image recognition as a model of the visual system, we show that such differences in representation can emerge as a direct consequence of different neural resource constraints on the retinal and cortical networks, and for the first time we find a single model from which both geometries spontaneously emerge at the appropriate stages of visual processing.

16, TITLE: Pay Less Attention with Lightweight and Dynamic Convolutions

<https://openreview.net/forum?id=SkVlh09tX>

AUTHORS: Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, Michael Auli

TL;DR: Dynamic lightweight convolutions are competitive to self-attention on language tasks.

HIGHLIGHT: In this paper, we show that a very lightweight convolution can perform competitively to the best reported self-attention results.

17, TITLE: Enabling Factorized Piano Music Modeling and Generation with the MAESTRO Dataset

<https://openreview.net/forum?id=r11YRjC9F7>

AUTHORS: Curtis Hawthorne, Andriy Stasyuk, Adam Roberts, Ian Simon, Cheng-Zhi Anna Huang, Sander Dieleman, Erich Elsen, Jesse Engel, Douglas Eck

TL;DR: We train a suite of models capable of transcribing, composing, and synthesizing audio waveforms with coherent musical structure, enabled by the new MAESTRO dataset.

HIGHLIGHT: The networks and the dataset together present a promising approach toward creating new expressive and interpretable neural models of music.

- 18, TITLE: Learning to Remember More with Less Memorization  
<https://openreview.net/forum?id=r1xlvi0qYm>  
AUTHORS: Hung Le, Truyen Tran, Svetha Venkatesh  
TL;DR: n/a  
HIGHLIGHT: This method aims to balance between maximizing memorization and forgetting via overwriting mechanisms.
- 19, TITLE: Learning Robust Representations by Projecting Superficial Statistics Out  
<https://openreview.net/forum?id=rJEjjoR9K7>  
AUTHORS: Haohan Wang, Zexue He, Zachary C. Lipton, Eric P. Xing  
TL;DR: Building on previous work on domain generalization, we hope to produce a classifier that will generalize to previously unseen domains, even when domain identifiers are not available during training.  
HIGHLIGHT: We test our method on the battery of standard domain generalization data sets and, interestingly, achieve comparable or better performance as compared to other domain generalization methods that explicitly require samples from the target distribution for training.
- 20, TITLE: Slalom: Fast, Verifiable and Private Execution of Neural Networks in Trusted Hardware  
<https://openreview.net/forum?id=rJVorjCcKQ>  
AUTHORS: Florian Tramèr, Dan Boneh  
TL;DR: We accelerate secure DNN inference in trusted execution environments (by a factor 4x-20x) by selectively outsourcing the computation of linear layers to a faster yet untrusted co-processor.  
HIGHLIGHT: Building upon an efficient outsourcing scheme for matrix multiplication, we propose Slalom, a framework that securely delegates execution of all linear layers in a DNN from a TEE (e.g., Intel SGX or Sanctum) to a faster, yet untrusted, co-located processor.
- 21, TITLE: The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision  
<https://openreview.net/forum?id=rJgMlhRctm>  
AUTHORS: Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, Jiajun Wu  
TL;DR: We propose the Neuro-Symbolic Concept Learner (NS-CL), a model that learns visual concepts, words, and semantic parsing of sentences without explicit supervision on any of them.  
HIGHLIGHT: We propose the Neuro-Symbolic Concept Learner (NS-CL), a model that learns visual concepts, words, and semantic parsing of sentences without explicit supervision on any of them; instead, our model learns by simply looking at images and reading paired questions and answers.
- 22, TITLE: The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks  
<https://openreview.net/forum?id=rJl-b3RcF7>  
AUTHORS: Jonathan Frankle, Michael Carbin  
TL;DR: Feedforward neural networks that can have weights pruned after training could have had the same weights pruned before training  
HIGHLIGHT: Neural network pruning techniques can reduce the parameter counts of trained networks by over 90%, decreasing storage requirements and improving computational performance of inference without compromising accuracy.
- 23, TITLE: FFJORD: Free-Form Continuous Dynamics for Scalable Reversible Generative Models  
<https://openreview.net/forum?id=rJxgknCcK7>  
AUTHORS: Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, David Duvenaud  
TL;DR: We use continuous time dynamics to define a generative model with exact likelihoods and efficient sampling that is parameterized by unrestricted neural networks.  
HIGHLIGHT: In this paper, we use Hutchinson's trace estimator to give a scalable unbiased estimate of the log-density.

24, TITLE: How Powerful are Graph Neural Networks?

<https://openreview.net/forum?id=ryGs6iA5Km>

AUTHORS: Keyulu Xu\*, Weihua Hu\*, Jure Leskovec, Stefanie Jegelka

TL;DR: We develop theoretical foundations for the expressive power of GNNs and design a provably most powerful GNN.

HIGHLIGHT: Here, we present a theoretical framework for analyzing the expressive power of GNNs to capture different graph structures.