1, TITLE: Efficient Algorithms for Non-convex Isotonic Regression through Submodular Optimization
https://papers.nips.cc/paper/7286-efficient-algorithms-for-non-convex-isotonic-regression-through-submodular-optimization
AUTHORS: Francis Bach
HIGHLIGHT: We propose new discretization schemes that lead to simple and efficient algorithms based on zero-th, first, or higher order oracles; these algorithms also lead to improvements without isotonic constraints.

2, TITLE: Structure-Aware Convolutional Neural Networks
https://papers.nips.cc/paper/7287-structure-aware-convolutional-neural-networks
AUTHORS: Jianlong Chang, Jie Gu, Lingfeng Wang, GAOFENG MENG, SHIMING XIANG, Chunhong Pan
HIGHLIGHT: To broaden the reach of CNNs, we develop structure-aware convolution to eliminate the invariance, yielding a unified mechanism of dealing with both Euclidean and non-Euclidean structured data.

3, TITLE: Kalman Normalization: Normalizing Internal Representations Across Network Layers
https://papers.nips.cc/paper/7288-kalman-normalization-normalizing-internal-representations-across-network-layers
AUTHORS: Guangrun Wang, jiefeng peng, Ping Luo, Xinjiang Wang, Liang Lin
HIGHLIGHT: In this paper, we present a novel normalization method, called Kalman Normalization (KN), for improving and accelerating the training of DNNs, particularly under the context of micro-batches.

4, TITLE: HOGWILD!-Gibbs can be PanAccurate
https://papers.nips.cc/paper/7289-hogwild-gibbs-can-be-panaccurate
AUTHORS: Constantinos Daskalakis, Nishanth Dikkala, Siddhartha Jayanti
HIGHLIGHT: We investigate whether it can be used to accurately estimate expectations of functions of {\em all the variables} of the model.

5, TITLE: Text-Adaptive Generative Adversarial Networks: Manipulating Images with Natural Language
https://papers.nips.cc/paper/7290-text-adaptive-generative-adversarial-networks-manipulating-images-with-natural-language
AUTHORS: Seonghyeon Nam, Yunji Kim, Seon Joo Kim
HIGHLIGHT: In this paper, we propose the text-adaptive generative adversarial network (TAGAN) to generate semantically manipulated images while preserving text-irrelevant contents.

6, TITLE: IntroVAE: Introspective Variational Autoencoders for Photographic Image Synthesis
https://papers.nips.cc/paper/7291-introvae-introspective-variational-autoencoders-for-photographic-image-synthesis
AUTHORS: Huaibo Huang, zhihang li, Ran He, Zhenan Sun, Tieniu Tan
HIGHLIGHT: We present a novel introspective variational autoencoder (IntroVAE) model for synthesizing high-resolution photographic images.

7, TITLE: Doubly Robust Bayesian Inference for Non-Stationary Streaming Data with <var>\beta</var>-Divergences
https://papers.nips.cc/paper/7292-doubly-robust-bayesian-inference-for-non-stationary-streaming-data-with-beta-divergences
AUTHORS: Jeremias Knoblauch, Jack E. Jewson, Theodoros Damoulas
HIGHLIGHT: We present the very first robust Bayesian Online Changepoint Detection algorithm through General Bayesian Inference (GBI) with β-divergences.

8, TITLE: Adapted Deep Embeddings: A Synthesis of Methods for k-Shot Inductive Transfer Learning
https://papers.nips.cc/paper/7293-adapted-deep-embeddings-a-synthesis-of-methods-for-k-shot-inductive-transfer-learning
AUTHORS: Tyler Scott, Karl Ridgeway, Michael C. Mozer
HIGHLIGHT: In weight transfer, a model trained on the source domain is used as an initialization point for a network to be trained on the target domain.

9, TITLE: Generalized Inverse Optimization through Online Learning
https://papers.nips.cc/paper/7294-generalized-inverse-optimization-through-online-learning
AUTHORS: Chaosheng Dong, Yiran Chen, Bo Zeng
HIGHLIGHT: In this paper, we propose a general framework for inverse optimization through online learning.

10, TITLE: An Off-policy Policy Gradient Theorem Using Emphatic Weightings
https://papers.nips.cc/paper/7295-an-off-policy-policy-gradient-theorem-using-emphatic-weightings
AUTHORS: Ehsan Imani, Eric Graves, Martha White
HIGHLIGHT: In this work, we solve this open problem by providing the first off-policy policy gradient theorem.

11, TITLE: Supervised autoencoders: Improving generalization performance with unsupervised regularizers
https://papers.nips.cc/paper/7296-supervised-autoencoders-improving-generalization-performance-with-unsupervised-regularizers

AUTHORS: Lei Le, Andrew Patterson, Martha White
HIGHLIGHT: In this work, we theoretically and empirically analyze one such model, called a supervised auto-encoder: a neural network that predicts both inputs (reconstruction error) and targets jointly.

12, TITLE: Visual Object Networks: Image Generation with Disentangled 3D Representations
https://papers.nips.cc/paper/7297-visual-object-networks-image-generation-with-disentangled-3d-representations
AUTHORS: Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, Bill Freeman
HIGHLIGHT: Inspired by classic graphics rendering pipelines, we unravel the image formation process into three conditionally independent factors---shape, viewpoint, and texture---and present an end-to-end adversarial learning framework that jointly models 3D shape and 2D texture.

13, TITLE: Understanding Weight Normalized Deep Neural Networks with Rectified Linear Units
https://papers.nips.cc/paper/7298-understanding-weight-normalized-deep-neural-networks-with-rectified-linear-units
AUTHORS: Yixi Xu, Xiao Wang
HIGHLIGHT: This paper presents a general framework for norm-based capacity control for $L_{p,q}$ weight normalized deep neural networks.

14, TITLE: Learning Pipelines with Limited Data and Domain Knowledge: A Study in Parsing Physics Problems
https://papers.nips.cc/paper/7299-learning-pipelines-with-limited-data-and-domain-knowledge-a-study-in-parsing-physics-problems
AUTHORS: Mrinmaya Sachan, Kumar Avinava Dubey, Tom M. Mitchell, Dan Roth, Eric P. Xing
HIGHLIGHT: As a case study, we present such a system that learns to parse Newtonian physics problems in textbooks.

15, TITLE: Learning long-range spatial dependencies with horizontal gated recurrent units
https://papers.nips.cc/paper/7300-learning-long-range-spatial-dependencies-with-horizontal-gated-recurrent-units
AUTHORS: Drew Linsley, Junkyung Kim, Vijay Veerabadran, Charles Windolf, Thomas Serre
HIGHLIGHT: We introduce a visual challenge, Pathfinder, and describe a novel recurrent neural network architecture called the horizontal gated recurrent unit (hGRU) to learn intrinsic horizontal connections -- both within and across feature columns.

16, TITLE: Joint Sub-bands Learning with Clique Structures for Wavelet Domain Super-Resolution
https://papers.nips.cc/paper/7301-joint-sub-bands-learning-with-clique-structures-for-wavelet-domain-super-resolution
AUTHORS: Zhisheng Zhong, Tiancheng Shen, Yibo Yang, Zhouchen Lin, Chao Zhang
HIGHLIGHT: To solve these problems, we propose the Super-Resolution CliqueNet (SRCliqueNet) to reconstruct the high resolution (HR) image with better textural details in the wavelet domain.

17, TITLE: Fast Similarity Search via Optimal Sparse Lifting
https://papers.nips.cc/paper/7302-fast-similarity-search-via-optimal-sparse-lifting
AUTHORS: Wenye Li, Jingwei Mao, Yin Zhang, Shuguang Cui
HIGHLIGHT: Motivated by the evidence in biological science, our work develops a novel approach for similarity search.

18, TITLE: Learning Deep Disentangled Embeddings With the F-Statistic Loss
https://papers.nips.cc/paper/7303-learning-deep-disentangled-embeddings-with-the-f-statistic-loss
AUTHORS: Karl Ridgeway, Michael C. Mozer
HIGHLIGHT: The goal of our work is to obtain more interpretable, manipulable, and generalizable deep representations of concepts and categories.

19, TITLE: Geometrically Coupled Monte Carlo Sampling
https://papers.nips.cc/paper/7304-geometrically-coupled-monte-carlo-sampling
AUTHORS: Mark Rowland, Krzysztof M. Choromanski, Fran?ois Chalus, Aldo Pacchiano, Tamas Sarlos, Richard E. Turner, Adrian Weller
HIGHLIGHT: We improve current methods for sampling in Euclidean spaces by avoiding independence, and instead consider ways to couple samples.

20, TITLE: Cooperative Holistic Scene Understanding: Unifying 3D Object, Layout, and Camera Pose Estimation
https://papers.nips.cc/paper/7305-cooperative-holistic-scene-understanding-unifying-3d-object-layout-and-camera-pose-estimation
AUTHORS: Siyuan Huang, Siyuan Qi, Yinxue Xiao, Yixin Zhu, Ying Nian Wu, Song-Chun Zhu
HIGHLIGHT: In this paper, we propose an end-to-end model that simultaneously solves all three tasks in real-time given only a single RGB image.

21, TITLE: An Efficient Pruning Algorithm for Robust Isotonic Regression
https://papers.nips.cc/paper/7306-an-efficient-pruning-algorithm-for-robust-isotonic-regression

AUTHORS: Cong Han Lim
HIGHLIGHT: We can combine techniques from the convex case with branch-and-bound ideas to form a new algorithm for this problem that naturally exploits the shape of the objective function.

22, TITLE: PAC-learning in the presence of adversaries
https://papers.nips.cc/paper/7307-pac-learning-in-the-presence-of-adversaries
AUTHORS: Daniel Cullina, Arjun Nitin Bhagoji, Prateek Mittal
HIGHLIGHT: In this paper, we step away from the attack-defense arms race and seek to understand the limits of what can be learned in the presence of an evasion adversary.

23, TITLE: Sparse DNNs with Improved Adversarial Robustness
https://papers.nips.cc/paper/7308-sparse-dnns-with-improved-adversarial-robustness
AUTHORS: Yiwen Guo, Chao Zhang, Changshui Zhang, Yurong Chen
HIGHLIGHT: Our analyses reveal, both theoretically and empirically, that nonlinear DNN-based classifiers behave differently under l_2 attacks from some linear ones.

24, TITLE: Snap ML: A Hierarchical Framework for Machine Learning
https://papers.nips.cc/paper/7309-snap-ml-a-hierarchical-framework-for-machine-learning
AUTHORS: Celestine D?nner, Thomas Parnell, Dimitrios Sarigiannis, Nikolas Ioannou, Andreea Anghel, Gummadi Ravi, Madhusudanan Kandasamy, Haralampos Pozidis
HIGHLIGHT: We describe a new software framework for fast training of generalized linear models.
Finally, we present a logistic regression benchmark on the Criteo Terabyte Click Logs dataset and show that Snap ML achieves the same test loss an order of magnitude faster than any of the previously reported results, including those obtained using TensorFlow and scikit-learn.

25, TITLE: See and Think: Disentangling Semantic Scene Completion
https://papers.nips.cc/paper/7310-see-and-think-disentangling-semantic-scene-completion
AUTHORS: Shice Liu, YU HU, Yiming Zeng, Qiankun Tang, Beibei Jin, Yinhe Han, Xiaowei Li
HIGHLIGHT: In this work, we propose a disentangled framework, sequentially carrying out 2D semantic segmentation, 2D-3D reprojection and 3D semantic scene completion.

26, TITLE: Chain of Reasoning for Visual Question Answering
https://papers.nips.cc/paper/7311-chain-of-reasoning-for-visual-question-answering
AUTHORS: Chenfei Wu, Jinlai Liu, Xiaojie Wang, Xuan Dong
HIGHLIGHT: This paper proposes a novel reasoning model for addressing these problems.

27, TITLE: Sigsoftmax: Reanalysis of the Softmax Bottleneck
https://papers.nips.cc/paper/7312-sigsoftmax-reanalysis-of-the-softmax-bottleneck
AUTHORS: Sekitoshi Kanai, Yasuhiro Fujiwara, Yuki Yamanaka, Shuichi Adachi
HIGHLIGHT: In this paper, we propose an output activation function for breaking the softmax bottleneck without additional parameters.

28, TITLE: Deep Non-Blind Deconvolution via Generalized Low-Rank Approximation
https://papers.nips.cc/paper/7313-deep-non-blind-deconvolution-via-generalized-low-rank-approximation
AUTHORS: Wenqi Ren, Jiawei Zhang, Lin Ma, Jinshan Pan, Xiaochun Cao, Wangmeng Zuo, Wei Liu, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we present a deep convolutional neural network to capture the inherent properties of image degradation, which can handle different kernels and saturated pixels in a unified framework.

29, TITLE: Bayesian Pose Graph Optimization via Bingham Distributions and Tempered Geodesic MCMC
https://papers.nips.cc/paper/7314-bayesian-pose-graph-optimization-via-bingham-distributions-and-tempered-geodesic-mcmc
AUTHORS: Tolga Birdal, Umut Simsekli, Mustafa Onur Eken, Slobodan Ilic
HIGHLIGHT: We introduce Tempered Geodesic Markov Chain Monte Carlo (TG-MCMC) algorithm for initializing pose graph optimization problems, arising in various scenarios such as SFM (structure from motion) or SLAM (simultaneous localization and mapping).

30, TITLE: MetaAnchor: Learning to Detect Objects with Customized Anchors
https://papers.nips.cc/paper/7315-metaanchor-learning-to-detect-objects-with-customized-anchors
AUTHORS: Tong Yang, Xiangyu Zhang, Zeming Li, Wenqiang Zhang, Jian Sun
HIGHLIGHT: We propose a novel and flexible anchor mechanism named MetaAnchor for object detection frameworks.

31, TITLE: Image Inpainting via Generative Multi-column Convolutional Neural Networks
https://papers.nips.cc/paper/7316-image-inpainting-via-generative-multi-column-convolutional-neural-networks
AUTHORS: Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, Jiaya Jia
HIGHLIGHT: In this paper, we propose a generative multi-column network for image inpainting.

32, TITLE: On Misinformation Containment in Online Social Networks
https://papers.nips.cc/paper/7317-on-misinformation-containment-in-online-social-networks
AUTHORS: Amo Tong, Ding-Zhu Du, Weili Wu
HIGHLIGHT: Motivated by realistic scenarios, we present the first analysis of the misinformation containment problem for the case when an arbitrary number of cascades are allowed.

33, TITLE: A^2-Nets: Double Attention Networks
https://papers.nips.cc/paper/7318-a2-nets-double-attention-networks
AUTHORS: Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, Jiashi Feng
HIGHLIGHT: In this work, we propose the “double attention block”, a novel component that aggregates and propagates informative global features from the entire spatio-temporal space of input images/videos, enabling subsequent convolution layers to access features from the entire space efficiently.

34, TITLE: Self-Supervised Generation of Spatial Audio for 360? Video
https://papers.nips.cc/paper/7319-self-supervised-generation-of-spatial-audio-for-360-video
AUTHORS: Pedro Morgado, Nuno Nvasconcelos, Timothy Langlois, Oliver Wang
HIGHLIGHT: We introduce an approach to convert mono audio recorded by a 360° video camera into spatial audio, a representation of the distribution of sound over the full viewing sphere.
We introduce several datasets, including one filmed ourselves, and one collected in-the-wild from YouTube, consisting of 360° videos uploaded with spatial audio.

35, TITLE: How Many Samples are Needed to Estimate a Convolutional Neural Network?
https://papers.nips.cc/paper/7320-how-many-samples-are-needed-to-estimate-a-convolutional-neural-network
AUTHORS: Simon S. Du, Yining Wang, Xiyu Zhai, Sivaraman Balakrishnan, Ruslan R. Salakhutdinov, Aarti Singh
HIGHLIGHT: For this model, we show that the sample complexity is $\widetilde{O}\left((m+r)/\epsilon^2\right)$ when the ratio between the stride size and the filter size is a constant.

36, TITLE: Algorithmic Regularization in Learning Deep Homogeneous Models: Layers are Automatically Balanced
https://papers.nips.cc/paper/7321-algorithmic-regularization-in-learning-deep-homogeneous-models-layers-are-automatically-balanced
AUTHORS: Simon S. Du, Wei Hu, Jason D. Lee
HIGHLIGHT: We study the implicit regularization imposed by gradient descent for learning multi-layer homogeneous functions including feed-forward fully connected and convolutional deep neural networks with linear, ReLU or Leaky ReLU activation.

37, TITLE: Optimization for Approximate Submodularity
https://papers.nips.cc/paper/7322-optimization-for-approximate-submodularity
AUTHORS: Yaron Singer, Avinatan Hassidim
HIGHLIGHT: We describe a technique which we call the sampled mean approximation that yields strong guarantees for maximization of submodular functions from approximate surrogates under cardinality and intersection of matroid constraints.

38, TITLE: (Probably) Concave Graph Matching
https://papers.nips.cc/paper/7323-probably-concave-graph-matching
AUTHORS: Haggai Maron, Yaron Lipman
HIGHLIGHT: In this paper we address the graph matching problem.

39, TITLE: Deep Defense: Training DNNs with Improved Adversarial Robustness
https://papers.nips.cc/paper/7324-deep-defense-training-dnns-with-improved-adversarial-robustness
AUTHORS: Ziang Yan, Yiwen Guo, Changshui Zhang
HIGHLIGHT: To address this problem, we propose a training recipe named "deep defense".

40, TITLE: Rest-Katyusha: Exploiting the Solution's Structure via Scheduled Restart Schemes
https://papers.nips.cc/paper/7325-rest-katyusha-exploiting-the-solutions-structure-via-scheduled-restart-schemes
AUTHORS: Junqi Tang, Mohammad Golbabaee, Francis Bach, Mike E. davies
HIGHLIGHT: We propose a structure-adaptive variant of the state-of-the-art stochastic variance-reduced gradient algorithm Katyusha for regularized empirical risk minimization.

41, TITLE: Implicit Reparameterization Gradients
https://papers.nips.cc/paper/7326-implicit-reparameterization-gradients
AUTHORS: Mikhail Figurnov, Shakir Mohamed, Andriy Mnih
HIGHLIGHT: We introduce an alternative approach to computing reparameterization gradients based on implicit differentiation and demonstrate its broader applicability by applying it to Gamma, Beta, Dirichlet, and von Mises distributions, which cannot be used with the classic reparameterization trick.

42, TITLE: Training DNNs with Hybrid Block Floating Point
https://papers.nips.cc/paper/7327-training-dnns-with-hybrid-block-floating-point
AUTHORS: Mario Drumond, Tao LIN, Martin Jaggi, Babak Falsafi
HIGHLIGHT: In this work, we introduce HBFP, a hybrid BFP-FP approach, which performs all dot products in BFP and other operations in floating point.

43, TITLE: A Model for Learned Bloom Filters and Optimizing by Sandwiching
https://papers.nips.cc/paper/7328-a-model-for-learned-bloom-filters-and-optimizing-by-sandwiching
AUTHORS: Michael Mitzenmacher
HIGHLIGHT: Here we model such learned Bloom filters, with the following outcomes: (1) we clarify what guarantees can and cannot be associated with such a structure; (2) we show how to estimate what size the learning function must obtain in order to obtain improved performance; (3) we provide a simple method, sandwiching, for optimizing learned Bloom filters; and (4) we propose a design and analysis approach for a learned Bloomier filter, based on our modeling approach.

44, TITLE: Soft-Gated Warping-GAN for Pose-Guided Person Image Synthesis
https://papers.nips.cc/paper/7329-soft-gated-warping-gan-for-pose-guided-person-image-synthesis
AUTHORS: Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, Jian Yin
HIGHLIGHT: This paper aims to resolve these challenges induced by geometric variability and spatial displacements via a new Soft-Gated Warping Generative Adversarial Network (Warping-GAN), which is composed of two stages: 1) it first synthesizes a target part segmentation map given a target pose, which depicts the region-level spatial layouts for guiding image synthesis with higher-level structure constraints; 2) the Warping-GAN equipped with a soft-gated warping-block learns feature-level mapping to render textures from the original image into the generated segmentation map.

45, TITLE: Deep Functional Dictionaries: Learning Consistent Semantic Structures on 3D Models from Functions
https://papers.nips.cc/paper/7330-deep-functional-dictionaries-learning-consistent-semantic-structures-on-3d-models-from-functions
AUTHORS: Minhyuk Sung, Hao Su, Ronald Yu, Leonidas J. Guibas
HIGHLIGHT: Given a collection of related 3D shapes, we consider how to jointly analyze such probe functions over different shapes, and how to discover common latent structures using a neural network — even in the absence of any correspondence information.

46, TITLE: Nonlocal Neural Networks, Nonlocal Diffusion and Nonlocal Modeling
https://papers.nips.cc/paper/7331-nonlocal-neural-networks-nonlocal-diffusion-and-nonlocal-modeling
AUTHORS: Yunzhe Tao, Qi Sun, Qiang Du, Wei Liu
HIGHLIGHT: Nonlocal Neural Networks, Nonlocal Diffusion and Nonlocal Modeling

47, TITLE: Are ResNets Provably Better than Linear Predictors?
https://papers.nips.cc/paper/7332-are-resnets-provably-better-than-linear-predictors
AUTHORS: Ohad Shamir
HIGHLIGHT: In this paper, we rigorously prove that arbitrarily deep, nonlinear residual units indeed exhibit this behavior, in the sense that the optimization landscape contains no local minima with value above what can be obtained with a linear predictor (namely a 1-layer network).

48, TITLE: Learning to Decompose and Disentangle Representations for Video Prediction
https://papers.nips.cc/paper/7333-learning-to-decompose-and-disentangle-representations-for-video-prediction
AUTHORS: Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li F. Fei-Fei, Juan Carlos Niebles
HIGHLIGHT: We address this challenge by proposing the Decompositional Disentangled Predictive Auto-Encoder (DDPAE), a framework that combines structured probabilistic models and deep networks to automatically (i) decompose the high-dimensional video that we aim to predict into components, and (ii) disentangle each component to have low-dimensional temporal dynamics that are easier to predict.

49, TITLE: Multi-Task Learning as Multi-Objective Optimization
https://papers.nips.cc/paper/7334-multi-task-learning-as-multi-objective-optimization
AUTHORS: Ozan Sener, Vladlen Koltun

HIGHLIGHT: In this paper, we explicitly cast multi-task learning as multi-objective optimization, with the overall objective of finding a Pareto optimal solution.

50, TITLE: Combinatorial Optimization with Graph Convolutional Networks and Guided Tree Search
https://papers.nips.cc/paper/7335-combinatorial-optimization-with-graph-convolutional-networks-and-guided-tree-search
AUTHORS: Zhuwen Li, Qifeng Chen, Vladlen Koltun
HIGHLIGHT: We present a learning-based approach to computing solutions for certain NP-hard problems.

51, TITLE: Self-Erasing Network for Integral Object Attention
https://papers.nips.cc/paper/7336-self-erasing-network-for-integral-object-attention
AUTHORS: Qibin Hou, PengTao Jiang, Yunchao Wei, Ming-Ming Cheng
HIGHLIGHT: To tackle such an issue as well as promote the quality of object attention, we introduce a simple yet effective Self-Erasing Network (SeeNet) to prohibit attentions from spreading to unexpected background regions.

52, TITLE: LinkNet: Relational Embedding for Scene Graph
https://papers.nips.cc/paper/7337-linknet-relational-embedding-for-scene-graph
AUTHORS: Sanghyun Woo, Dahun Kim, Donghyeon Cho, In So Kweon
HIGHLIGHT: In this paper, we present a novel method that improves scene graph generation by explicitly modeling inter-dependency among the entire object instances.

53, TITLE: How to Start Training: The Effect of Initialization and Architecture
https://papers.nips.cc/paper/7338-how-to-start-training-the-effect-of-initialization-and-architecture
AUTHORS: Boris Hanin, David Rolnick
HIGHLIGHT: We identify and study two common failure modes for early training in deep ReLU nets.

54, TITLE: Which Neural Net Architectures Give Rise to Exploding and Vanishing Gradients?
https://papers.nips.cc/paper/7339-which-neural-net-architectures-give-rise-to-exploding-and-vanishing-gradients
AUTHORS: Boris Hanin
HIGHLIGHT: We give a rigorous analysis of the statistical behavior of gradients in a randomly initialized fully connected network N with ReLU activations.

55, TITLE: Explanations based on the Missing: Towards Contrastive Explanations with Pertinent Negatives
https://papers.nips.cc/paper/7340-explanations-based-on-the-missing-towards-contrastive-explanations-with-pertinent-negatives
AUTHORS: Amit Dhurandhar, Pin-Yu Chen, Ronny Luss, Chun-Chen Tu, Paishun Ting, Karthikeyan Shanmugam, Payel Das
HIGHLIGHT: In this paper we propose a novel method that provides contrastive explanations justifying the classification of an input by a black box classifier such as a deep neural network.

56, TITLE: HitNet: Hybrid Ternary Recurrent Neural Network
https://papers.nips.cc/paper/7341-hitnet-hybrid-ternary-recurrent-neural-network
AUTHORS: Peiqi Wang, Xinfeng Xie, Lei Deng, Guoqi Li, Dongsheng Wang, Yuan Xie
HIGHLIGHT: Based on our observation, we propose HitNet, a hybrid ternary recurrent neural network, which bridges the accuracy gap between the full precision model and the quantized model.

57, TITLE: A Unified Framework for Extensive-Form Game Abstraction with Bounds
https://papers.nips.cc/paper/7342-a-unified-framework-for-extensive-form-game-abstraction-with-bounds
AUTHORS: Christian Kroer, Tuomas Sandholm
HIGHLIGHT: In this paper we present a unified framework for analyzing abstractions that can express all types of abstractions and solution concepts used in prior papers with performance guarantees---while maintaining comparable bounds on abstraction quality.

58, TITLE: Removing the Feature Correlation Effect of Multiplicative Noise
https://papers.nips.cc/paper/7343-removing-the-feature-correlation-effect-of-multiplicative-noise
AUTHORS: Zijun Zhang, Yining Zhang, Zongpeng Li
HIGHLIGHT: In this work, we propose non-correlating multiplicative noise (NCMN), which exploits batch normalization to remove the correlation effect in a simple yet effective way.

59, TITLE: Maximum-Entropy Fine Grained Classification
https://papers.nips.cc/paper/7344-maximum-entropy-fine-grained-classification
AUTHORS: Abhimanyu Dubey, Otkrist Gupta, Ramesh Raskar, Nikhil Naik

HIGHLIGHT: Utilizing this notion of small visual diversity, we revisit Maximum-Entropy learning in the context of fine-grained classification, and provide a training routine that maximizes the entropy of the output probability distribution for training convolutional neural networks on FGVC tasks.

60, TITLE: On Learning Markov Chains
https://papers.nips.cc/paper/7345-on-learning-markov-chains
AUTHORS: Yi HAO, Alon Orlitsky, Venkatadheeraj Pichapati
HIGHLIGHT: We consider two problems related to the min-max risk (expected loss) of estimating an unknown k-state Markov chain from its n sequential samples: predicting the conditional distribution of the next sample with respect to the KL-divergence, and estimating the transition matrix with respect to a natural loss induced by KL or a more general f-divergence measure.

61, TITLE: A Neural Compositional Paradigm for Image Captioning
https://papers.nips.cc/paper/7346-a-neural-compositional-paradigm-for-image-captioning
AUTHORS: Bo Dai, Sanja Fidler, Dahua Lin
HIGHLIGHT: In this paper, we present an alternative paradigm for image captioning, which factorizes the captioning procedure into two stages: (1) extracting an explicit semantic representation from the given image; and (2) constructing the caption based on a recursive compositional procedure in a bottom-up manner.

62, TITLE: Quantifying Learning Guarantees for Convex but Inconsistent Surrogates
https://papers.nips.cc/paper/7347-quantifying-learning-guarantees-for-convex-but-inconsistent-surrogates
AUTHORS: Kirill Struminsky, Simon Lacoste-Julien, Anton Osokin
HIGHLIGHT: Our key technical contribution consists in a new lower bound on the calibration function for the quadratic surrogate, which is non-trivial (not always zero) for inconsistent cases.

63, TITLE: Dialog-based Interactive Image Retrieval
https://papers.nips.cc/paper/7348-dialog-based-interactive-image-retrieval
AUTHORS: Xiaoxiao Guo, Hui Wu, Yu Cheng, Steven Rennie, Gerald Tesauro, Rogerio Feris
HIGHLIGHT: In this paper, we introduce a new approach to interactive image search that enables users to provide feedback via natural language, allowing for more natural and effective interaction.

64, TITLE: SPIDER: Near-Optimal Non-Convex Optimization via Stochastic Path-Integrated Differential Estimator
https://papers.nips.cc/paper/7349-spider-near-optimal-non-convex-optimization-via-stochastic-path-integrated-differential-estimator
AUTHORS: Cong Fang, Chris Junchi Li, Zhouchen Lin, Tong Zhang
HIGHLIGHT: In this paper, we propose a new technique named \textit{Stochastic Path-Integrated Differential EstimatoR} (SPIDER), which can be used to track many deterministic quantities of interests with significantly reduced computational cost.

65, TITLE: Are GANs Created Equal? A Large-Scale Study
https://papers.nips.cc/paper/7350-are-gans-created-equal-a-large-scale-study
AUTHORS: Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, Olivier Bousquet
HIGHLIGHT: We conduct a neutral, multi-faceted large-scale empirical study on state-of-the art models and evaluation measures.
To overcome some limitations of the current metrics, we also propose several data sets on which precision and recall can be computed.

66, TITLE: Learning Disentangled Joint Continuous and Discrete Representations
https://papers.nips.cc/paper/7351-learning-disentangled-joint-continuous-and-discrete-representations
AUTHORS: Emilien Dupont
HIGHLIGHT: We present a framework for learning disentangled and interpretable jointly continuous and discrete representations in an unsupervised manner.

67, TITLE: TADAM: Task dependent adaptive metric for improved few-shot learning
https://papers.nips.cc/paper/7352-tadam-task-dependent-adaptive-metric-for-improved-few-shot-learning
AUTHORS: Boris Oreshkin, Pau Rodr?guez L?pez, Alexandre Lacoste
HIGHLIGHT: In this work, we identify that metric scaling and metric task conditioning are important to improve the performance of few-shot algorithms.

68, TITLE: Do Less, Get More: Streaming Submodular Maximization with Subsampling
https://papers.nips.cc/paper/7353-do-less-get-more-streaming-submodular-maximization-with-subsampling
AUTHORS: Moran Feldman, Amin Karbasi, Ehsan Kazemi
HIGHLIGHT: In this paper, we develop the first one-pass streaming algorithm for submodular maximization that does not evaluate the entire stream even once.

69, TITLE: Deep Neural Nets with Interpolating Function as Output Activation
https://papers.nips.cc/paper/7355-deep-neural-nets-with-interpolating-function-as-output-activation
AUTHORS: Bao Wang, Xiyang Luo, Zhen Li, Wei Zhu, Zuoqiang Shi, Stanley Osher
HIGHLIGHT: And we propose end-to-end training and testing algorithms for this new architecture.

70, TITLE: FishNet: A Versatile Backbone for Image, Region, and Pixel Level Prediction
https://papers.nips.cc/paper/7356-fishnet-a-versatile-backbone-for-image-region-and-pixel-level-prediction
AUTHORS: Shuyang Sun, Jiangmiao Pang, Jianping Shi, Shuai Yi, Wanli Ouyang
HIGHLIGHT: Towards this goal, we design a fish-like network, called FishNet.

71, TITLE: Visual Memory for Robust Path Following
https://papers.nips.cc/paper/7357-visual-memory-for-robust-path-following
AUTHORS: Ashish Kumar, Saurabh Gupta, David Fouhey, Sergey Levine, Jitendra Malik
HIGHLIGHT: In this paper, we present an approach for doing so.

72, TITLE: KDGAN: Knowledge Distillation with Generative Adversarial Networks
https://papers.nips.cc/paper/7358-kdgan-knowledge-distillation-with-generative-adversarial-networks
AUTHORS: Xiaojie Wang, Rui Zhang, Yu Sun, Jianzhong Qi
HIGHLIGHT: To address these limitations, we propose a three-player game named KDGAN consisting of a classifier, a teacher, and a discriminator.

73, TITLE: Long short-term memory and Learning-to-learn in networks of spiking neurons
https://papers.nips.cc/paper/7359-long-short-term-memory-and-learning-to-learn-in-networks-of-spiking-neurons
AUTHORS: Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, Wolfgang Maass
HIGHLIGHT: We address two possible reasons for that.

74, TITLE: Greedy Hash: Towards Fast Optimization for Accurate Hash Coding in CNN
https://papers.nips.cc/paper/7360-greedy-hash-towards-fast-optimization-for-accurate-hash-coding-in-cnn
AUTHORS: Shupeng Su, Chao Zhang, Kai Han, Yonghong Tian
HIGHLIGHT: In this work, we adopt the greedy principle to tackle this NP hard problem by iteratively updating the network toward the probable optimal discrete solution in each iteration.

75, TITLE: Informative Features for Model Comparison
https://papers.nips.cc/paper/7361-informative-features-for-model-comparison
AUTHORS: Wittawat Jitkrittum, Heishiro Kanagawa, Patsorn Sangkloy, James Hays, Bernhard Sch?lkopf, Arthur Gretton
HIGHLIGHT: We propose two new statistical tests which are nonparametric, computationally efficient (runtime complexity is linear in the sample size), and interpretable.

76, TITLE: PointCNN: Convolution On X-Transformed Points
https://papers.nips.cc/paper/7362-pointcnn-convolution-on-x-transformed-points
AUTHORS: Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, Baoquan Chen
HIGHLIGHT: We present a simple and general framework for feature learning from point cloud.

77, TITLE: Connectionist Temporal Classification with Maximum Entropy Regularization
https://papers.nips.cc/paper/7363-connectionist-temporal-classification-with-maximum-entropy-regularization
AUTHORS: Hu Liu, Sheng Jin, Changshui Zhang
HIGHLIGHT: To remedy this, we propose a regularization method based on maximum conditional entropy which penalizes peaky distributions and encourages exploration.

78, TITLE: Large Margin Deep Networks for Classification
https://papers.nips.cc/paper/7364-large-margin-deep-networks-for-classification
AUTHORS: Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, Samy Bengio
HIGHLIGHT: In this work, we propose a novel loss function to impose a margin on any chosen set of layers of a deep network (including input and hidden layers).

79, TITLE: Generalizing Graph Matching beyond Quadratic Assignment Model
https://papers.nips.cc/paper/7365-generalizing-graph-matching-beyond-quadratic-assignment-model
AUTHORS: Tianshu Yu, Junchi Yan, Yilin Wang, Wei Liu, baoxin Li

HIGHLIGHT: We show that a large family of functions, which we define as Separable Functions, can approximate discrete graph matching in the continuous domain asymptotically by varying the approximation controlling parameters.

80, TITLE: Solving Large Sequential Games with the Excessive Gap Technique
https://papers.nips.cc/paper/7366-solving-large-sequential-games-with-the-excessive-gap-technique
AUTHORS: Christian Kroer, Gabriele Farina, Tuomas Sandholm
HIGHLIGHT: In this paper we show that a particular first-order method, a state-of-the-art variant of the excessive gap technique---instantiated with the dilated entropy distance function---can efficiently solve large real-world problems competitively with CFR and its variants.

81, TITLE: Discrimination-aware Channel Pruning for Deep Neural Networks
https://papers.nips.cc/paper/7367-discrimination-aware-channel-pruning-for-deep-neural-networks
AUTHORS: Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, Jinhui Zhu
HIGHLIGHT: To this end, we introduce additional losses into the network to increase the discriminative power of intermediate layers and then select the most discriminative channels for each layer by considering the additional loss and the reconstruction error.

82, TITLE: On the Dimensionality of Word Embedding
https://papers.nips.cc/paper/7368-on-the-dimensionality-of-word-embedding
AUTHORS: Zi Yin, Yuanyuan Shen
HIGHLIGHT: In this paper, we provide a theoretical understanding of word embedding and its dimensionality.

83, TITLE: Reinforced Continual Learning
https://papers.nips.cc/paper/7369-reinforced-continual-learning
AUTHORS: Ju Xu, Zhanxing Zhu
HIGHLIGHT: In this work, a novel approach for continual learning is proposed, which searches for the best neural architecture for each coming task via sophisticatedly designed reinforcement learning strategies.

84, TITLE: Uncertainty-Aware Attention for Reliable Interpretation and Prediction
https://papers.nips.cc/paper/7370-uncertainty-aware-attention-for-reliable-interpretation-and-prediction
AUTHORS: Jay Heo, Hae Beom Lee, Saehoon Kim, Juho Lee, Kwang Joon Kim, Eunho Yang, Sung Ju Hwang
HIGHLIGHT: To overcome this limitation, we introduce the notion of input-dependent uncertainty to the attention mechanism, such that it generates attention for each feature with varying degrees of noise based on the given input, to learn larger variance on instances it is uncertain about.

85, TITLE: DropMax: Adaptive Variational Softmax
https://papers.nips.cc/paper/7371-dropmax-adaptive-variational-softmax
AUTHORS: Hae Beom Lee, Juho Lee, Saehoon Kim, Eunho Yang, Sung Ju Hwang
HIGHLIGHT: We propose DropMax, a stochastic version of softmax classifier which at each iteration drops non-target classes according to dropout probabilities adaptively decided for each instance.

86, TITLE: Posterior Concentration for Sparse Deep Learning
https://papers.nips.cc/paper/7372-posterior-concentration-for-sparse-deep-learning
AUTHORS: Veronika Rockova, nicholas polson
HIGHLIGHT: We introduce Spike-and-Slab Deep Learning (SS-DL), a fully Bayesian alternative to dropout for improving generalizability of deep ReLU networks.

87, TITLE: A flexible model for training action localization with varying levels of supervision
https://papers.nips.cc/paper/7373-a-flexible-model-for-training-action-localization-with-varying-levels-of-supervision
AUTHORS: Guilhem Ch?ron, Jean-Baptiste Alayrac, Ivan Laptev, Cordelia Schmid
HIGHLIGHT: In this work we propose a unifying framework that can handle and combine varying types of less demanding weak supervision.

88, TITLE: A Deep Bayesian Policy Reuse Approach Against Non-Stationary Agents
https://papers.nips.cc/paper/7374-a-deep-bayesian-policy-reuse-approach-against-non-stationary-agents
AUTHORS: YAN ZHENG, Zhaopeng Meng, Jianye Hao, Zongzhang Zhang, Tianpei Yang, Changjie Fan
HIGHLIGHT: To detect policy accurately, we propose the \textit{rectified belief model} taking advantage of the \textit{opponent model} to infer the other agent's policy from reward signals and its behaviors.

89, TITLE: Empirical Risk Minimization in Non-interactive Local Differential Privacy Revisited

https://papers.nips.cc/paper/7375-empirical-risk-minimization-in-non-interactive-local-differential-privacy-revisited
AUTHORS: Di Wang, Marco Gaboardi, Jinhui Xu
HIGHLIGHT: In this paper, we revisit the Empirical Risk Minimization problem in the non-interactive local model of differential privacy.

90, TITLE: Low-shot Learning via Covariance-Preserving Adversarial Augmentation Networks
https://papers.nips.cc/paper/7376-low-shot-learning-via-covariance-preserving-adversarial-augmentation-networks
AUTHORS: Hang Gao, Zheng Shou, Alireza Zareian, Hanwang Zhang, Shih-Fu Chang
HIGHLIGHT: In this work, we propose Covariance-Preserving Adversarial Augmentation Networks to overcome existing limits of low-shot learning.

91, TITLE: Learning semantic similarity in a continuous space
https://papers.nips.cc/paper/7377-learning-semantic-similarity-in-a-continuous-space
AUTHORS: Michel Deudon
HIGHLIGHT: We address the problem of learning semantic representation of questions to measure similarity between pairs as a continuous distance metric.

92, TITLE: MetaReg: Towards Domain Generalization using Meta-Regularization
https://papers.nips.cc/paper/7378-metareg-towards-domain-generalization-using-meta-regularization
AUTHORS: Yogesh Balaji, Swami Sankaranarayanan, Rama Chellappa
HIGHLIGHT: In this work, we encode this notion of domain generalization using a novel regularization function.
We pose the problem of finding such a regularization function in a Learning to Learn (or) meta-learning framework.

93, TITLE: Boosted Sparse and Low-Rank Tensor Regression
https://papers.nips.cc/paper/7379-boosted-sparse-and-low-rank-tensor-regression
AUTHORS: Lifang He, Kun Chen, Wanwan Xu, Jiayu Zhou, Fei Wang
HIGHLIGHT: We propose a sparse and low-rank tensor regression model to relate a univariate outcome to a feature tensor, in which each unit-rank tensor from the CP decomposition of the coefficient tensor is assumed to be sparse.

94, TITLE: Domain-Invariant Projection Learning for Zero-Shot Recognition
https://papers.nips.cc/paper/7380-domain-invariant-projection-learning-for-zero-shot-recognition
AUTHORS: An Zhao, Mingyu Ding, Jiechao Guan, Zhiwu Lu, Tao Xiang, Ji-Rong Wen
HIGHLIGHT: In this paper, we propose a novel ZSL model termed domain-invariant projection learning (DIPL).

95, TITLE: Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding
https://papers.nips.cc/paper/7381-neural-symbolic-vqa-disentangling-reasoning-from-vision-and-language-understanding
AUTHORS: Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, Josh Tenenbaum
HIGHLIGHT: Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding

96, TITLE: Frequency-Domain Dynamic Pruning for Convolutional Neural Networks
https://papers.nips.cc/paper/7382-frequency-domain-dynamic-pruning-for-convolutional-neural-networks
AUTHORS: Zhenhua Liu, Jizheng Xu, Xiulian Peng, Ruiqin Xiong
HIGHLIGHT: Considering that there are spatial redundancy within most filters in a CNN, we propose a frequency-domain dynamic pruning scheme to exploit the spatial correlations.

97, TITLE: Quadratic Decomposable Submodular Function Minimization
https://papers.nips.cc/paper/7383-quadratic-decomposable-submodular-function-minimization
AUTHORS: Pan Li, Niao He, Olgica Milenkovic
HIGHLIGHT: We approach the problem via a new dual strategy and describe an objective that may be optimized via random coordinate descent (RCD) methods and projections onto cones.
We introduce a new convex optimization problem, termed quadratic decomposable submodular function minimization.

98, TITLE: A Block Coordinate Ascent Algorithm for Mean-Variance Optimization
https://papers.nips.cc/paper/7384-a-block-coordinate-ascent-algorithm-for-mean-variance-optimization
AUTHORS: Tengyang Xie, Bo Liu, Yangyang Xu, Mohammad Ghavamzadeh, Yinlam Chow, Daoming Lyu, Daesub Yoon
HIGHLIGHT: In this paper, we develop a model-free policy search framework for mean-variance optimization with finite-sample error bound analysis (to local optima).

99, TITLE: <var>\ell_1</var>-regression with Heavy-tailed Distributions
https://papers.nips.cc/paper/7385-ell_1-regression-with-heavy-tailed-distributions

AUTHORS: Lijun Zhang, Zhi-Hua Zhou
HIGHLIGHT: In this paper, we consider the problem of linear regression with heavy-tailed distributions.
To address the challenge that both the input and output could be heavy-tailed, we propose a truncated minimization problem, and demonstrate that it enjoys an $O(\sqrt{d/n})$ excess risk, where d is the dimensionality and n is the number of samples.

100, TITLE: Neural Nearest Neighbors Networks
https://papers.nips.cc/paper/7386-neural-nearest-neighbors-networks
AUTHORS: Tobias Pl?tz, Stefan Roth
HIGHLIGHT: To overcome this, we propose a continuous deterministic relaxation of KNN selection that maintains differentiability w.r.t. pairwise distances, but retains the original KNN as the limit of a temperature parameter approaching zero.

101, TITLE: Efficient nonmyopic batch active search
https://papers.nips.cc/paper/7387-efficient-nonmyopic-batch-active-search
AUTHORS: Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, Roman Garnett
HIGHLIGHT: We bridge this gap, addressing batch active search from both the theoretical and practical perspective.

102, TITLE: A Game-Theoretic Approach to Recommendation Systems with Strategic Content Providers
https://papers.nips.cc/paper/7388-a-game-theoretic-approach-to-recommendation-systems-with-strategic-content-providers
AUTHORS: Omer Ben-Porat, Moshe Tennenholtz
HIGHLIGHT: We introduce a game-theoretic approach to the study of recommendation systems with strategic content providers.

103, TITLE: Interactive Structure Learning with Structural Query-by-Committee
https://papers.nips.cc/paper/7389-interactive-structure-learning-with-structural-query-by-committee
AUTHORS: Christopher Tosh, Sanjoy Dasgupta
HIGHLIGHT: In this work, we introduce interactive structure learning, a framework that unifies many different interactive learning tasks.

104, TITLE: Global Geometry of Multichannel Sparse Blind Deconvolution on the Sphere
https://papers.nips.cc/paper/7390-global-geometry-of-multichannel-sparse-blind-deconvolution-on-the-sphere
AUTHORS: Yanjun Li, Yoram Bresler
HIGHLIGHT: We consider the case where the x_i's are sparse, and convolution with f is invertible.

105, TITLE: Video-to-Video Synthesis
https://papers.nips.cc/paper/7391-video-to-video-synthesis
AUTHORS: Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, Bryan Catanzaro
HIGHLIGHT: In this paper, we propose a video-to-video synthesis approach under the generative adversarial learning framework.

106, TITLE: How To Make the Gradients Small Stochastically: Even Faster Convex and Nonconvex SGD
https://papers.nips.cc/paper/7392-how-to-make-the-gradients-small-stochastically-even-faster-convex-and-nonconvex-sgd
AUTHORS: Zeyuan Allen-Zhu
HIGHLIGHT: If $f(x)$ is convex, to find a point with gradient norm ε, we design an algorithm SGD3 with a near-optimal rate $\tilde{O}(\varepsilon^{-2})$, improving the best known rate $O(\varepsilon^{-8/3})$.

107, TITLE: Synthesized Policies for Transfer and Adaptation across Tasks and Environments
https://papers.nips.cc/paper/7393-synthesized-policies-for-transfer-and-adaptation-across-tasks-and-environments
AUTHORS: Hexiang Hu, Liyu Chen, Boqing Gong, Fei Sha
HIGHLIGHT: In this paper, we consider the problem of learning to simultaneously transfer across both environments and tasks, probably more importantly, by learning from only sparse (environment, task) pairs out of all the possible combinations.

108, TITLE: Adversarial vulnerability for any classifier
https://papers.nips.cc/paper/7394-adversarial-vulnerability-for-any-classifier
AUTHORS: Alhussein Fawzi, Hamza Fawzi, Omar Fawzi
HIGHLIGHT: In this paper, we study the phenomenon of adversarial perturbations under the assumption that the data is generated with a smooth generative model.

109, TITLE: Evolution-Guided Policy Gradient in Reinforcement Learning
https://papers.nips.cc/paper/7395-evolution-guided-policy-gradient-in-reinforcement-learning
AUTHORS: Shauharda Khadka, Kagan Tumer

HIGHLIGHT: In this paper, we introduce Evolutionary Reinforcement Learning (ERL), a hybrid algorithm that leverages the population of an EA to provide diversified data to train an RL agent, and reinserts the RL agent into the EA population periodically to inject gradient information into the EA.

110, TITLE: Toddler-Inspired Visual Object Learning
https://papers.nips.cc/paper/7396-toddler-inspired-visual-object-learning
AUTHORS: Sven Bambach, David Crandall, Linda Smith, Chen Yu
HIGHLIGHT: Using head-mounted cameras, eye gaze trackers, and a model of foveated vision, we collected first-person (egocentric) images that represents a highly accurate approximation of the "training data" that toddlers' visual systems collect in everyday, naturalistic learning contexts.

111, TITLE: Alternating optimization of decision trees, with application to learning sparse oblique trees
https://papers.nips.cc/paper/7397-alternating-optimization-of-decision-trees-with-application-to-learning-sparse-oblique-trees
AUTHORS: Miguel A. Carreira-Perpinan, Pooya Tavallali
HIGHLIGHT: We give an algorithm that, given an input tree (its structure and the parameter values at its nodes), produces a new tree with the same or smaller structure but new parameter values that provably lower or leave unchanged the misclassification error.

112, TITLE: FD-GAN: Pose-guided Feature Distilling GAN for Robust Person Re-identification
https://papers.nips.cc/paper/7398-fd-gan-pose-guided-feature-distilling-gan-for-robust-person-re-identification
AUTHORS: Yixiao Ge, Zhuowan Li, Haiyu Zhao, Guojun Yin, Shuai Yi, Xiaogang Wang, hongsheng Li
HIGHLIGHT: FD-GAN: Pose-guided Feature Distilling GAN for Robust Person Re-identification

113, TITLE: New Insight into Hybrid Stochastic Gradient Descent: Beyond With-Replacement Sampling and Convexity
https://papers.nips.cc/paper/7399-new-insight-into-hybrid-stochastic-gradient-descent-beyond-with-replacement-sampling-and-convexity
AUTHORS: Pan Zhou, Xiaotong Yuan, Jiashi Feng
HIGHLIGHT: In this paper, we affirmatively answer this open question by showing that under WoRS and for both convex and non-convex problems, it is still possible for HSGD (with constant step-size) to match full gradient descent in rate of convergence, while maintaining comparable sample-size-independent incremental first-order oracle complexity to stochastic gradient descent.

114, TITLE: The Lingering of Gradients: How to Reuse Gradients Over Time
https://papers.nips.cc/paper/7400-the-lingering-of-gradients-how-to-reuse-gradients-over-time
AUTHORS: Zeyuan Allen-Zhu, David Simchi-Levi, Xinshang Wang
HIGHLIGHT: In this paper, we study a more refined complexity by taking into account the ``lingering'' of gradients: once a gradient is computed at x_k, the additional time to compute gradients at $x_{k+1},x_{k+2},\dots$ may be reduced.

115, TITLE: Unsupervised Learning of View-invariant Action Representations
https://papers.nips.cc/paper/7401-unsupervised-learning-of-view-invariant-action-representations
AUTHORS: Junnan Li, Yongkang Wong, Qi Zhao, Mohan Kankanhalli
HIGHLIGHT: In this work, we propose an unsupervised learning framework, which exploits unlabeled data to learn video representations.

116, TITLE: Fairness Behind a Veil of Ignorance: A Welfare Analysis for Automated Decision Making
https://papers.nips.cc/paper/7402-fairness-behind-a-veil-of-ignorance-a-welfare-analysis-for-automated-decision-making
AUTHORS: Hoda Heidari, Claudio Ferrari, Krishna Gummadi, Andreas Krause
HIGHLIGHT: We draw attention to an important, yet largely overlooked aspect of evaluating fairness for automated decision making systems---namely risk and welfare considerations.

117, TITLE: Global Gated Mixture of Second-order Pooling for Improving Deep Convolutional Neural Networks
https://papers.nips.cc/paper/7403-global-gated-mixture-of-second-order-pooling-for-improving-deep-convolutional-neural-networks
AUTHORS: Qilong Wang, Zilin Gao, Jiangtao Xie, Wangmeng Zuo, Peihua Li
HIGHLIGHT: To this end, we introduce a sparsity-constrained gating mechanism and propose a novel parametric SOP as component of mixture model.

118, TITLE: Image-to-image translation for cross-domain disentanglement
https://papers.nips.cc/paper/7404-image-to-image-translation-for-cross-domain-disentanglement
AUTHORS: Abel Gonzalez-Garcia, Joost van de Weijer, Yoshua Bengio
HIGHLIGHT: In this paper, we bridge these two objectives and introduce the concept of cross-domain disentanglement.

119, TITLE: Gradient Sparsification for Communication-Efficient Distributed Optimization
https://papers.nips.cc/paper/7405-gradient-sparsification-for-communication-efficient-distributed-optimization
AUTHORS: Jianqiao Wangni, Jialei Wang, Ji Liu, Tong Zhang
HIGHLIGHT: In this paper, to reduce the communication cost, we propose a convex optimization formulation to minimize the coding length of stochastic gradients.

120, TITLE: Revisiting Multi-Task Learning with ROCK: a Deep Residual Auxiliary Block for Visual Detection
https://papers.nips.cc/paper/7406-revisiting-multi-task-learning-with-rock-a-deep-residual-auxiliary-block-for-visual-detection
AUTHORS: Taylor Mordan, Nicolas THOME, Gilles Henaff, Matthieu Cord
HIGHLIGHT: In this paper, we tackle a specific MTL context denoted as primary MTL, where the ultimate goal is to improve the performance of a given primary task by leveraging several other auxiliary tasks.

121, TITLE: Adaptive Online Learning in Dynamic Environments
https://papers.nips.cc/paper/7407-adaptive-online-learning-in-dynamic-environments
AUTHORS: Lijun Zhang, Shiyin Lu, Zhi-Hua Zhou
HIGHLIGHT: In this paper, we study online convex optimization in dynamic environments, and aim to bound the dynamic regret with respect to any sequence of comparators.

122, TITLE: FRAGE: Frequency-Agnostic Word Representation
https://papers.nips.cc/paper/7408-frage-frequency-agnostic-word-representation
AUTHORS: Chengyue Gong, Di He, Xu Tan, Tao Qin, Liwei Wang, Tie-Yan Liu
HIGHLIGHT: In order to mitigate the issue, in this paper, we propose a neat, simple yet effective adversarial training method to blur the boundary between the embeddings of high-frequency words and low-frequency words.

123, TITLE: Generative Neural Machine Translation
https://papers.nips.cc/paper/7409-generative-neural-machine-translation
AUTHORS: Harshil Shah, David Barber
HIGHLIGHT: We introduce Generative Neural Machine Translation (GNMT), a latent variable architecture which is designed to model the semantics of the source and target sentences.

124, TITLE: Found Graph Data and Planted Vertex Covers
https://papers.nips.cc/paper/7410-found-graph-data-and-planted-vertex-covers
AUTHORS: Austin R. Benson, Jon Kleinberg
HIGHLIGHT: We develop a framework for analyzing this planted vertex cover problem, based on the theory of fixed-parameter tractability, together with algorithms for recovering the core.

125, TITLE: Joint Active Feature Acquisition and Classification with Variable-Size Set Encoding
https://papers.nips.cc/paper/7411-joint-active-feature-acquisition-and-classification-with-variable-size-set-encoding
AUTHORS: Hajin Shim, Sung Ju Hwang, Eunho Yang
HIGHLIGHT: In this work, we formulate this active feature acquisition as a jointly learning problem of training both the classifier (environment) and the RL agent that decides either to `stop and predict' or `collect a new feature' at test time, in a cost-sensitive manner.

126, TITLE: Regularization Learning Networks: Deep Learning for Tabular Datasets
https://papers.nips.cc/paper/7412-regularization-learning-networks-deep-learning-for-tabular-datasets
AUTHORS: Ira Shavitt, Eran Segal
HIGHLIGHT: We propose that applying a different regularization coefficient to each weight might boost the performance of DNNs by allowing them to make more use of the more relevant inputs.

127, TITLE: Multitask Boosting for Survival Analysis with Competing Risks
https://papers.nips.cc/paper/7413-multitask-boosting-for-survival-analysis-with-competing-risks
AUTHORS: Alexis Bellot, Mihaela van der Schaar
HIGHLIGHT: We introduce in this paper a survival model with the flexibility to leverage a common representation of related events that is designed to correct for the strong imbalance in observed outcomes.

128, TITLE: Geometry Based Data Generation
https://papers.nips.cc/paper/7414-geometry-based-data-generation
AUTHORS: Ofir Lindenbaum, Jay Stanley, Guy Wolf, Smita Krishnaswamy
HIGHLIGHT: We propose a new type of generative model for high-dimensional data that learns a manifold geometry of the data, rather than density, and can generate points evenly along this manifold.

129, TITLE: SLAYER: Spike Layer Error Reassignment in Time
https://papers.nips.cc/paper/7415-slayer-spike-layer-error-reassignment-in-time
AUTHORS: Sumit Bam Shrestha, Garrick Orchard
HIGHLIGHT: In this paper, we introduce a new general backpropagation mechanism for learning synaptic weights and axonal delays which overcomes the problem of non-differentiability of the spike function and uses a temporal credit assignment policy for backpropagating error to preceding layers.

130, TITLE: On Oracle-Efficient PAC RL with Rich Observations
https://papers.nips.cc/paper/7416-on-oracle-efficient-pac-rl-with-rich-observations
AUTHORS: Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert E. Schapire
HIGHLIGHT: We present new provably sample-efficient algorithms for environments with deterministic hidden state dynamics and stochastic rich observations.

131, TITLE: Gradient Descent for Spiking Neural Networks
https://papers.nips.cc/paper/7417-gradient-descent-for-spiking-neural-networks
AUTHORS: Dongsung Huh, Terrence J. Sejnowski
HIGHLIGHT: Here, we present a gradient descent method for optimizing spiking network models by introducing a differentiable formulation of spiking dynamics and deriving the exact gradient calculation.

132, TITLE: Generalizing Tree Probability Estimation via Bayesian Networks
https://papers.nips.cc/paper/7418-generalizing-tree-probability-estimation-via-bayesian-networks
AUTHORS: Cheng Zhang, Frederick A Matsen IV
HIGHLIGHT: In this paper, we derive a general Bayesian network formulation for probability estimation on leaf-labeled trees that enables flexible approximations which can generalize beyond observations.

133, TITLE: Where Do You Think You're Going?: Inferring Beliefs about Dynamics from Behavior
https://papers.nips.cc/paper/7419-where-do-you-think-youre-going-inferring-beliefs-about-dynamics-from-behavior
AUTHORS: Sid Reddy, Anca Dragan, Sergey Levine
HIGHLIGHT: In this paper, we take an alternative approach, and model suboptimal behavior as the result of internal model misspecification: the reason that user actions might deviate from near-optimal actions is that the user has an incorrect set of beliefs about the rules -- the dynamics -- governing how actions affect the environment.

134, TITLE: Designing by Training: Acceleration Neural Network for Fast High-Dimensional Convolution
https://papers.nips.cc/paper/7420-designing-by-training-acceleration-neural-network-for-fast-high-dimensional-convolution
AUTHORS: Longquan Dai, Liang Tang, Yuan Xie, Jinhui Tang
HIGHLIGHT: Instead, we propose an Acceleration Network (AccNet) which turns the work of designing new fast algorithms to training the AccNet.

135, TITLE: Understanding the Role of Adaptivity in Machine Teaching: The Case of Version Space Learners
https://papers.nips.cc/paper/7421-understanding-the-role-of-adaptivity-in-machine-teaching-the-case-of-version-space-learners
AUTHORS: Yuxin Chen, Adish Singla, Oisin Mac Aodha, Pietro Perona, Yisong Yue
HIGHLIGHT: In this paper, we study the case of teaching consistent, version space learners in an interactive setting.

136, TITLE: A loss framework for calibrated anomaly detection
https://papers.nips.cc/paper/7422-a-loss-framework-for-calibrated-anomaly-detection
AUTHORS:
HIGHLIGHT: This paper concerns calibrated anomaly detection, which is the practically relevant extension where we additionally wish to produce a confidence score for a point being anomalous.

137, TITLE: PacGAN: The power of two samples in generative adversarial networks
https://papers.nips.cc/paper/7423-pacgan-the-power-of-two-samples-in-generative-adversarial-networks
AUTHORS: Zinan Lin, Ashish Khetan, Giulia Fanti, Sewoong Oh
HIGHLIGHT: We study a principled approach to handling mode collapse, which we call packing.

138, TITLE: Variational Memory Encoder-Decoder
https://papers.nips.cc/paper/7424-variational-memory-encoder-decoder
AUTHORS: Hung Le, Truyen Tran, Thin Nguyen, Svetha Venkatesh
HIGHLIGHT: To overcome this, we explore a novel approach that injects variability into neural encoder-decoder via the use of external memory as a mixture model, namely Variational Memory Encoder-Decoder (VMED).

139, TITLE: Stochastic Composite Mirror Descent: Optimal Bounds with High Probabilities
https://papers.nips.cc/paper/7425-stochastic-composite-mirror-descent-optimal-bounds-with-high-probabilities
AUTHORS: Yunwen Lei, Ke Tang
HIGHLIGHT: We study stochastic composite mirror descent, a class of scalable algorithms able to exploit the geometry and composite structure of a problem.

140, TITLE: Hybrid Retrieval-Generation Reinforced Agent for Medical Image Report Generation
https://papers.nips.cc/paper/7426-hybrid-retrieval-generation-reinforced-agent-for-medical-image-report-generation
AUTHORS: Yuan Li, Xiaodan Liang, Zhiting Hu, Eric P. Xing
HIGHLIGHT: Generating long and coherent reports to describe medical images poses challenges to bridging visual patterns with informative human linguistic descriptions.

141, TITLE: Overcoming Language Priors in Visual Question Answering with Adversarial Regularization
https://papers.nips.cc/paper/7427-overcoming-language-priors-in-visual-question-answering-with-adversarial-regularization
AUTHORS: Sainandan Ramakrishnan, Aishwarya Agrawal, Stefan Lee
HIGHLIGHT: In this work, we present a novel regularization scheme for VQA that reduces this effect.

142, TITLE: Hybrid Knowledge Routed Modules for Large-scale Object Detection
https://papers.nips.cc/paper/7428-hybrid-knowledge-routed-modules-for-large-scale-object-detection
AUTHORS: ChenHan Jiang, Hang Xu, Xiaodan Liang, Liang Lin
HIGHLIGHT: Particularly, we present Hybrid Knowledge Routed Modules (HKRM) that incorporates the reasoning routed by two kinds of knowledge forms: an explicit knowledge module for structured constraints that are summarized with linguistic knowledge (e.g. shared attributes, relationships) about concepts; and an implicit knowledge module that depicts some implicit constraints (e.g. common spatial layouts).

143, TITLE: Bilinear Attention Networks
https://papers.nips.cc/paper/7429-bilinear-attention-networks
AUTHORS: Jin-Hwa Kim, Jaehyun Jun, Byoung-Tak Zhang
HIGHLIGHT: In this paper, we propose bilinear attention networks (BAN) that find bilinear attention distributions to utilize given vision-language information seamlessly.

144, TITLE: Parsimonious Quantile Regression of Financial Asset Tail Dynamics via Sequential Learning
https://papers.nips.cc/paper/7430-parsimonious-quantile-regression-of-financial-asset-tail-dynamics-via-sequential-learning
AUTHORS: Xing Yan, Weizhong Zhang, Lin Ma, Wei Liu, Qi Wu
HIGHLIGHT: We propose a parsimonious quantile regression framework to learn the dynamic tail behaviors of financial asset returns.

145, TITLE: Multi-Class Learning: From Theory to Algorithm
https://papers.nips.cc/paper/7431-multi-class-learning-from-theory-to-algorithm
AUTHORS: Jian Li, Yong Liu, Rong Yin, Hua Zhang, Lizhong Ding, Weiping Wang
HIGHLIGHT: In this paper, we study the generalization performance of multi-class classification and obtain a shaper data-dependent generalization error bound with fast convergence rate, substantially improving the state-of-art bounds in the existing data-dependent generalization analysis.

146, TITLE: Multivariate Time Series Imputation with Generative Adversarial Networks
https://papers.nips.cc/paper/7432-multivariate-time-series-imputation-with-generative-adversarial-networks
AUTHORS: Yonghong Luo, Xiangrui Cai, Ying ZHANG, Jun Xu, Yuan xiaojie
HIGHLIGHT: In this paper, we treat the problem of missing value imputation as data generation.

147, TITLE: Learning Versatile Filters for Efficient Convolutional Neural Networks
https://papers.nips.cc/paper/7433-learning-versatile-filters-for-efficient-convolutional-neural-networks
AUTHORS: Yunhe Wang, Chang Xu, Chunjing XU, Chao Xu, Dacheng Tao
HIGHLIGHT: This paper introduces versatile filters to construct efficient convolutional neural network.

148, TITLE: Accelerated Stochastic Matrix Inversion: General Theory and Speeding up BFGS Rules for Faster Second-Order Optimization
https://papers.nips.cc/paper/7434-accelerated-stochastic-matrix-inversion-general-theory-and-speeding-up-bfgs-rules-for-faster-second-order-optimization
AUTHORS: Robert Gower, Filip Hanzely, Peter Richtarik, Sebastian U. Stich
HIGHLIGHT: We present the first accelerated randomized algorithm for solving linear systems in Euclidean spaces.

149, TITLE: DifNet: Semantic Segmentation by Diffusion Networks
https://papers.nips.cc/paper/7435-difnet-semantic-segmentation-by-diffusion-networks
AUTHORS: Peng Jiang, Fanglin Gu, Yunhai Wang, Changhe Tu, Baoquan Chen
HIGHLIGHT: Instead, in this work, we decompose this difficult task into two relative simple sub-tasks: seed detection which is required to predict initial predictions without the need of wholeness and preciseness, and similarity estimation which measures the possibility of any two nodes belong to the same class without the need of knowing which class they are.

150, TITLE: Conditional Adversarial Domain Adaptation
https://papers.nips.cc/paper/7436-conditional-adversarial-domain-adaptation
AUTHORS: Mingsheng Long, ZHANGJIE CAO, Jianmin Wang, Michael I. Jordan
HIGHLIGHT: In this paper, we present conditional adversarial domain adaptation, a principled framework that conditions the adversarial adaptation models on discriminative information conveyed in the classifier predictions.

151, TITLE: Neighbourhood Consensus Networks
https://papers.nips.cc/paper/7437-neighbourhood-consensus-networks
AUTHORS: Ignacio Rocco, Mircea Cimpoi, Relja Arandjelovic, Akihiko Torii, Tomas Pajdla, Josef Sivic
HIGHLIGHT: The contributions of this work are threefold.

152, TITLE: Relating Leverage Scores and Density using Regularized Christoffel Functions
https://papers.nips.cc/paper/7438-relating-leverage-scores-and-density-using-regularized-christoffel-functions
AUTHORS: Edouard Pauwels, Francis Bach, Jean-Philippe Vert
HIGHLIGHT: Borrowing ideas from the orthogonal polynomial literature, we introduce the regularized Christoffel function associated to a positive definite kernel.

153, TITLE: Non-Local Recurrent Network for Image Restoration
https://papers.nips.cc/paper/7439-non-local-recurrent-network-for-image-restoration
AUTHORS: Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, Thomas S. Huang
HIGHLIGHT: The main contributions of this work are: (1) Unlike existing methods that measure self-similarity in an isolated manner, the proposed non-local module can be flexibly integrated into existing deep networks for end-to-end training to capture deep feature correlation between each location and its neighborhood.

154, TITLE: Bayesian Semi-supervised Learning with Graph Gaussian Processes
https://papers.nips.cc/paper/7440-bayesian-semi-supervised-learning-with-graph-gaussian-processes
AUTHORS: Yin Cheng Ng, Nicol? Colombo, Ricardo Silva
HIGHLIGHT: We propose a data-efficient Gaussian process-based Bayesian approach to the semi-supervised learning problem on graphs.

155, TITLE: Foreground Clustering for Joint Segmentation and Localization in Videos and Images
https://papers.nips.cc/paper/7441-foreground-clustering-for-joint-segmentation-and-localization-in-videos-and-images
AUTHORS: Abhishek Sharma
HIGHLIGHT: This paper presents a novel framework in which video/image segmentation and localization are cast into a single optimization problem that integrates information from low level appearance cues with that of high level localization cues in a very weakly supervised manner.

156, TITLE: Video Prediction via Selective Sampling
https://papers.nips.cc/paper/7442-video-prediction-via-selective-sampling
AUTHORS: Jingwei Xu, Bingbing Ni, Xiaokang Yang
HIGHLIGHT: Combining above two insights we propose a two-stage network called VPSS (\textbf{V}ideo \textbf{P}rediction via \textbf{S}elective \textbf{S}ampling).

157, TITLE: Distilled Wasserstein Learning for Word Embedding and Topic Modeling
https://papers.nips.cc/paper/7443-distilled-wasserstein-learning-for-word-embedding-and-topic-modeling
AUTHORS: Hongteng Xu, Wenlin Wang, Wei Liu, Lawrence Carin
HIGHLIGHT: We propose a novel Wasserstein method with a distillation mechanism, yielding joint learning of word embeddings and topics.

158, TITLE: Learning to Exploit Stability for 3D Scene Parsing
https://papers.nips.cc/paper/7444-learning-to-exploit-stability-for-3d-scene-parsing
AUTHORS: Yilun Du, Zhijian Liu, Hector Basevi, Ales Leonardis, Bill Freeman, Josh Tenenbaum, Jiajun Wu

HIGHLIGHT: We integrate the physical cue of stability into the learning process using a REINFORCE approach coupled to a physics engine, and apply this to the problem of producing the 3D bounding boxes and poses of objects in a scene.

159, TITLE: Neural Guided Constraint Logic Programming for Program Synthesis
https://papers.nips.cc/paper/7445-neural-guided-constraint-logic-programming-for-program-synthesis
AUTHORS: Lisa Zhang, Gregory Rosenblatt, Ethan Fetaya, Renjie Liao, William Byrd, Matthew Might, Raquel Urtasun, Richard Zemel
HIGHLIGHT: We present a method for solving Programming By Example (PBE) problems by using a neural model to guide the search of a constraint logic programming system called miniKanren.

160, TITLE: Genetic-Gated Networks for Deep Reinforcement Learning
https://papers.nips.cc/paper/7446-genetic-gated-networks-for-deep-reinforcement-learning
AUTHORS: Simyung Chang, John Yang, Jaeseok Choi, Nojun Kwak
HIGHLIGHT: We introduce the Genetic-Gated Networks (G2Ns), simple neural networks that combine a gate vector composed of binary genetic genes in the hidden layer(s) of networks.

161, TITLE: Fighting Boredom in Recommender Systems with Linear Reinforcement Learning
https://papers.nips.cc/paper/7447-fighting-boredom-in-recommender-systems-with-linear-reinforcement-learning
AUTHORS: Romain WARLOP, Alessandro Lazaric, J?r?mie Mary
HIGHLIGHT: We argue that this assumption is rarely verified in practice, as the recommendation process itself may impact the user's preferences.

162, TITLE: Enhancing the Accuracy and Fairness of Human Decision Making
https://papers.nips.cc/paper/7448-enhancing-the-accuracy-and-fairness-of-human-decision-making
AUTHORS: Isabel Valera, Adish Singla, Manuel Gomez Rodriguez
HIGHLIGHT: In this paper, we address the above problem from the perspective of sequential decision making and show that, for different fairness notions from the literature, it reduces to a sequence of (constrained) weighted bipartite matchings, which can be solved efficiently using algorithms with approximation guarantees.

163, TITLE: Temporal Regularization for Markov Decision Process
https://papers.nips.cc/paper/7449-temporal-regularization-for-markov-decision-process
AUTHORS: Pierre Thodoroff, Audrey Durand, Joelle Pineau, Doina Precup
HIGHLIGHT: Most existing regularization techniques focus on spatial (perceptual) regularization.

164, TITLE: The Pessimistic Limits and Possibilities of Margin-based Losses in Semi-supervised Learning
https://papers.nips.cc/paper/7450-the-pessimistic-limits-and-possibilities-of-margin-based-losses-in-semi-supervised-learning
AUTHORS: Jesse Krijthe, Marco Loog
HIGHLIGHT: We show that for linear classifiers defined by convex margin-based surrogate losses that are decreasing, it is impossible to construct \emph{any} semi-supervised approach that is able to guarantee an improvement over the supervised classifier measured by this surrogate loss on the labeled and unlabeled data.

165, TITLE: Simple random search of static linear policies is competitive for reinforcement learning
https://papers.nips.cc/paper/7451-simple-random-search-of-static-linear-policies-is-competitive-for-reinforcement-learning
AUTHORS: Horia Mania, Aurelia Guy, Benjamin Recht
HIGHLIGHT: We introduce a model-free random search algorithm for training static, linear policies for continuous control problems.

166, TITLE: Generating Informative and Diverse Conversational Responses via Adversarial Information Maximization
https://papers.nips.cc/paper/7452-generating-informative-and-diverse-conversational-responses-via-adversarial-information-maximization
AUTHORS: Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, Bill Dolan
HIGHLIGHT: We present Adversarial Information Maximization (AIM), an adversarial learning framework that addresses these two related but distinct problems.

167, TITLE: Entropy and mutual information in models of deep neural networks
https://papers.nips.cc/paper/7453-entropy-and-mutual-information-in-models-of-deep-neural-networks
AUTHORS: Marylou Gabri?, Andre Manoel, Cl?ment Luneau, jean barbier, Nicolas Macris, Florent Krzakala, Lenka Zdeborov?
HIGHLIGHT: (iii) We propose an experiment framework with generative models of synthetic datasets, on which we train deep neural networks with a weight constraint designed so that the assumption in (i) is verified during learning.

168, TITLE: Collaborative Learning for Deep Neural Networks
https://papers.nips.cc/paper/7454-collaborative-learning-for-deep-neural-networks
AUTHORS: Guocong Song, Wei Chai
HIGHLIGHT: We introduce collaborative learning in which multiple classifier heads of the same network are simultaneously trained on the same training data to improve generalization and robustness to label noise with no extra inference cost.

169, TITLE: High Dimensional Linear Regression using Lattice Basis Reduction
https://papers.nips.cc/paper/7455-high-dimensional-linear-regression-using-lattice-basis-reduction
AUTHORS: Ilias Zadik, David Gamarnik
HIGHLIGHT: We propose a new polynomial-time algorithm for this task which is based on the seminal Lenstra-Lenstra-Lovasz (LLL) lattice basis reduction algorithm.

170, TITLE: Symbolic Graph Reasoning Meets Convolutions
https://papers.nips.cc/paper/7456-symbolic-graph-reasoning-meets-convolutions
AUTHORS: Xiaodan Liang, Zhiting Hu, Hao Zhang, Liang Lin, Eric P. Xing
HIGHLIGHT: Rather than using separate graphical models (e.g. CRF) or constraints for modeling broader dependencies, we propose a new Symbolic Graph Reasoning (SGR) layer, which performs reasoning over a group of symbolic nodes whose outputs explicitly represent different properties of each semantic in a prior knowledge graph.

171, TITLE: DVAE#: Discrete Variational Autoencoders with Relaxed Boltzmann Priors
https://papers.nips.cc/paper/7457-dvae-discrete-variational-autoencoders-with-relaxed-boltzmann-priors
AUTHORS: Arash Vahdat, Evgeny Andriyash, William Macready
HIGHLIGHT: We propose two approaches for relaxing Boltzmann machines to continuous distributions that permit training with importance-weighted bounds.

172, TITLE: Partially-Supervised Image Captioning
https://papers.nips.cc/paper/7458-partially-supervised-image-captioning
AUTHORS: Peter Anderson, Stephen Gould, Mark Johnson
HIGHLIGHT: Since image labels and object classes can be interpreted as partial captions, we formulate this problem as learning from partially-specified sequence data.

173, TITLE: 3D-Aware Scene Manipulation via Inverse Graphics
https://papers.nips.cc/paper/7459-3d-aware-scene-manipulation-via-inverse-graphics
AUTHORS: Shunyu Yao, Tzu Ming Hsu, Jun-Yan Zhu, Jiajun Wu, Antonio Torralba, Bill Freeman, Josh Tenenbaum
HIGHLIGHT: In this work, we propose 3D scene de-rendering networks (3D-SDN) to address the above issues by integrating disentangled representations for semantics, geometry, and appearance into a deep generative model.

174, TITLE: Random Feature Stein Discrepancies
https://papers.nips.cc/paper/7460-random-feature-stein-discrepancies
AUTHORS: Jonathan Huggins, Lester Mackey
HIGHLIGHT: To address these shortcomings, we introduce feature Stein discrepancies (?SDs), a new family of quality measures that can be cheaply approximated using importance sampling.

175, TITLE: Distributed Stochastic Optimization via Adaptive SGD
https://papers.nips.cc/paper/7461-distributed-stochastic-optimization-via-adaptive-sgd
AUTHORS: Ashok Cutkosky, R?bert Busa-Fekete
HIGHLIGHT: In this paper, we propose an efficient distributed stochastic optimization method by combining adaptivity with variance reduction techniques.

176, TITLE: Precision and Recall for Time Series
https://papers.nips.cc/paper/7462-precision-and-recall-for-time-series
AUTHORS: Nesime Tatbul, Tae Jun Lee, Stan Zdonik, Mejbah Alam, Justin Gottschlich
HIGHLIGHT: Motivated by this observation, we present a new mathematical model to evaluate the accuracy of time series classification algorithms.

177, TITLE: Deep Attentive Tracking via Reciprocative Learning
https://papers.nips.cc/paper/7463-deep-attentive-tracking-via-reciprocative-learning
AUTHORS: Shi Pu, Yibing Song, Chao Ma, Honggang Zhang, Ming-Hsuan Yang
HIGHLIGHT: In this paper, we propose a reciprocative learning algorithm to exploit visual attention for training deep classifiers.

178, TITLE: Virtual Class Enhanced Discriminative Embedding Learning
https://papers.nips.cc/paper/7464-virtual-class-enhanced-discriminative-embedding-learning
AUTHORS: Binghui Chen, Weihong Deng, Haifeng Shen
HIGHLIGHT: In this paper, we propose a novel yet extremely simple method Virtual Softmax to enhance the discriminative property of learned features by injecting a dynamic virtual negative class into the original softmax.

179, TITLE: Attention in Convolutional LSTM for Gesture Recognition
https://papers.nips.cc/paper/7465-attention-in-convolutional-lstm-for-gesture-recognition
AUTHORS: Liang Zhang, Guangming Zhu, Lin Mei, Peiyi Shen, Syed Afaq Ali Shah, Mohammed Bennamoun
HIGHLIGHT: Based on the previous gesture recognition architectures which combine the three-dimensional convolution neural network (3DCNN) and ConvLSTM, this paper explores the effects of attention mechanism in ConvLSTM.

180, TITLE: Pelee: A Real-Time Object Detection System on Mobile Devices
https://papers.nips.cc/paper/7466-pelee-a-real-time-object-detection-system-on-mobile-devices
AUTHORS: Robert J. Wang, Xiang Li, Charles X. Ling
HIGHLIGHT: In this study, we propose an efficient architecture named PeleeNet, which is built with conventional convolution instead.

181, TITLE: Universal Growth in Production Economies
https://papers.nips.cc/paper/7467-universal-growth-in-production-economies
AUTHORS: Simina Branzei, Ruta Mehta, Noam Nisan
HIGHLIGHT: We study a simple variant of the von Neumann model of an expanding economy, in which multiple producers make goods according to their production function.

182, TITLE: Bayesian Model Selection Approach to Boundary Detection with Non-Local Priors
https://papers.nips.cc/paper/7468-bayesian-model-selection-approach-to-boundary-detection-with-non-local-priors
AUTHORS: Fei Jiang, Guosheng Yin, Francesca Dominici
HIGHLIGHT: Based on non-local prior distributions, we propose a Bayesian model selection (BMS) procedure for boundary detection in a sequence of data with multiple systematic mean changes.

183, TITLE: Efficient Stochastic Gradient Hard Thresholding
https://papers.nips.cc/paper/7469-efficient-stochastic-gradient-hard-thresholding
AUTHORS: Pan Zhou, Xiaotong Yuan, Jiashi Feng
HIGHLIGHT: To address these deficiencies, we propose an efficient hybrid stochastic gradient hard thresholding (HSG-HT) method that can be provably shown to have sample-size-independent gradient evaluation and hard thresholding complexity bounds.

184, TITLE: SplineNets: Continuous Neural Decision Graphs
https://papers.nips.cc/paper/7470-splinenets-continuous-neural-decision-graphs
AUTHORS: Cem Keskin, Shahram Izadi
HIGHLIGHT: We present SplineNets, a practical and novel approach for using conditioning in convolutional neural networks (CNNs).

185, TITLE: Generalized Zero-Shot Learning with Deep Calibration Network
https://papers.nips.cc/paper/7471-generalized-zero-shot-learning-with-deep-calibration-network
AUTHORS: Shichen Liu, Mingsheng Long, Jianmin Wang, Michael I. Jordan
HIGHLIGHT: In this paper, we study generalized zero-shot learning that assumes accessible to target classes for unseen data during training, and prediction on unseen data is made by searching on both source and target classes.

186, TITLE: Neural Architecture Search with Bayesian Optimisation and Optimal Transport
https://papers.nips.cc/paper/7472-neural-architecture-search-with-bayesian-optimisation-and-optimal-transport
AUTHORS: Kirthevasan Kandasamy, Willie Neiswanger, Jeff Schneider, Barnabas Poczos, Eric P. Xing
HIGHLIGHT: In this work, we develop NASBOT, a Gaussian process based BO framework for neural architecture search.

187, TITLE: Embedding Logical Queries on Knowledge Graphs
https://papers.nips.cc/paper/7473-embedding-logical-queries-on-knowledge-graphs
AUTHORS: Will Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, Jure Leskovec
HIGHLIGHT: Here we introduce a framework to efficiently make predictions about conjunctive logical queries -- a flexible but tractable subset of first-order logic -- on incomplete knowledge graphs.

188, TITLE: Learning Optimal Reserve Price against Non-myopic Bidders
https://papers.nips.cc/paper/7474-learning-optimal-reserve-price-against-non-myopic-bidders
AUTHORS: Jinyan Liu, Zhiyi Huang, Xiangning Wang
HIGHLIGHT: We introduce algorithms that obtain small regret against non-myopic bidders either when the market is large, i.e., no bidder appears in a constant fraction of the rounds, or when the bidders are impatient, i.e., they discount future utility by some factor mildly bounded away from one.

189, TITLE: Sequential Context Encoding for Duplicate Removal
https://papers.nips.cc/paper/7475-sequential-context-encoding-for-duplicate-removal
AUTHORS: Lu Qi, Shu Liu, Jianping Shi, Jiaya Jia
HIGHLIGHT: In this work, we design a new two-stage framework to effectively select the appropriate proposal candidate for each object.

190, TITLE: Discovery of Latent 3D Keypoints via End-to-end Geometric Reasoning
https://papers.nips.cc/paper/7476-discovery-of-latent-3d-keypoints-via-end-to-end-geometric-reasoning
AUTHORS: Supasorn Suwajanakorn, Noah Snavely, Jonathan J. Tompson, Mohammad Norouzi
HIGHLIGHT: This paper presents KeypointNet, an end-to-end geometric reasoning framework to learn an optimal set of category-specific keypoints, along with their detectors to predict 3D keypoints in a single 2D input image.

191, TITLE: Nonparametric learning from Bayesian models with randomized objective functions
https://papers.nips.cc/paper/7477-nonparametric-learning-from-bayesian-models-with-randomized-objective-functions
AUTHORS: Simon Lyddon, Stephen Walker, Chris C. Holmes
HIGHLIGHT: Here we present a Bayesian nonparametric approach to learning that makes use of statistical models, but does not assume that the model is true.

192, TITLE: SEGA: Variance Reduction via Gradient Sketching
https://papers.nips.cc/paper/7478-sega-variance-reduction-via-gradient-sketching
AUTHORS: Filip Hanzely, Konstantin Mishchenko, Peter Richtarik
HIGHLIGHT: We propose a novel randomized first order optimization method---SEGA (SkEtched GrAdient method)---which progressively throughout its iterations builds a variance-reduced estimate of the gradient from random linear measurements (sketches) of the gradient provided at each iteration by an oracle.

193, TITLE: Automatic Program Synthesis of Long Programs with a Learned Garbage Collector
https://papers.nips.cc/paper/7479-automatic-program-synthesis-of-long-programs-with-a-learned-garbage-collector
AUTHORS: Amit Zohar, Lior Wolf
HIGHLIGHT: We consider the problem of generating automatic code given sample input-output pairs.

194, TITLE: One-Shot Unsupervised Cross Domain Translation
https://papers.nips.cc/paper/7480-one-shot-unsupervised-cross-domain-translation
AUTHORS: Sagie Benaim, Lior Wolf
HIGHLIGHT: We argue that this task could be a key AI capability that underlines the ability of cognitive agents to act in the world and present empirical evidence that the existing unsupervised domain translation methods fail on this task.

195, TITLE: Regularizing by the Variance of the Activations' Sample-Variances
https://papers.nips.cc/paper/7481-regularizing-by-the-variance-of-the-activations-sample-variances
AUTHORS: Etai Littwin, Lior Wolf
HIGHLIGHT: As we prove, this encourages the activations to be distributed around a few distinct modes.

196, TITLE: Overlapping Clustering Models, and One (class) SVM to Bind Them All
https://papers.nips.cc/paper/7482-overlapping-clustering-models-and-one-class-svm-to-bind-them-all
AUTHORS: Xueyu Mao, Purnamrita Sarkar, Deepayan Chakrabarti
HIGHLIGHT: Many existing overlapping clustering methods model each person (or word, or book) as a non-negative weighted combination of "exemplars" who belong solely to one community, with some small noise.

197, TITLE: Algorithmic Linearly Constrained Gaussian Processes
https://papers.nips.cc/paper/7483-algorithmic-linearly-constrained-gaussian-processes
AUTHORS: Markus Lange-Hegermann
HIGHLIGHT: We algorithmically construct multi-output Gaussian process priors which satisfy linear differential equations.

198, TITLE: DeepExposure: Learning to Expose Photos with Asynchronously Reinforced Adversarial Learning
https://papers.nips.cc/paper/7484-deepexposure-learning-to-expose-photos-with-asynchronously-reinforced-adversarial-learning
AUTHORS: Runsheng Yu, Wenyu Liu, Yasen Zhang, Zhi Qu, Deli Zhao, Bo Zhang
HIGHLIGHT: Inspired by luminosity masks usually applied by professional photographers, in this paper, we develop a novel algorithm for learning local exposures with deep reinforcement adversarial learning.

199, TITLE: Norm matters: efficient and accurate normalization schemes in deep networks
https://papers.nips.cc/paper/7485-norm-matters-efficient-and-accurate-normalization-schemes-in-deep-networks
AUTHORS: Elad Hoffer, Ron Banner, Itay Golan, Daniel Soudry
HIGHLIGHT: In this work, we present a novel view on the purpose and function of normalization methods and weight-decay, as tools to decouple weights' norm from the underlying optimized objective.

200, TITLE: Dual Principal Component Pursuit: Improved Analysis and Efficient Algorithms
https://papers.nips.cc/paper/7486-dual-principal-component-pursuit-improved-analysis-and-efficient-algorithms
AUTHORS: Zhihui Zhu, Yifan Wang, Daniel Robinson, Daniel Naiman, Rene Vidal, Manolis Tsakiris
HIGHLIGHT: In this paper we provide a refined geometric analysis and a new statistical analysis that show that DPCP can tolerate as many outliers as the square of the number of inliers, thus improving upon other provably correct robust PCA methods.

201, TITLE: MULAN: A Blind and Off-Grid Method for Multichannel Echo Retrieval
https://papers.nips.cc/paper/7487-mulan-a-blind-and-off-grid-method-for-multichannel-echo-retrieval
AUTHORS: Helena Peic Tukuljac, Antoine Deleforge, Remi Gribonval
HIGHLIGHT: We propose a radically different approach to the problem, building on top of the framework of finite-rate-of-innovation sampling.

202, TITLE: Mixture Matrix Completion
https://papers.nips.cc/paper/7488-mixture-matrix-completion
AUTHORS: Daniel Pimentel-Alarcon
HIGHLIGHT: This paper generalizes these models to what we call mixture matrix completion (MMC): the case where each entry of X corresponds to one of several low-rank matrices.

203, TITLE: Trajectory Convolution for Action Recognition
https://papers.nips.cc/paper/7489-trajectory-convolution-for-action-recognition
AUTHORS: Yue Zhao, Yuanjun Xiong, Dahua Lin
HIGHLIGHT: In this work, we propose a new CNN architecture TrajectoryNet, which incorporates trajectory convolution, a new operation for integrating features along the temporal dimension, to replace the existing temporal convolution.

204, TITLE: The Description Length of Deep Learning models
https://papers.nips.cc/paper/7490-the-description-length-of-deep-learning-models
AUTHORS: L?onard Blier, Yann Ollivier
HIGHLIGHT: In this work, we show experimentally that despite their huge number of parameters, deep neural networks can compress the data losslessly even when taking the cost of encoding the parameters into account.

205, TITLE: A Smoothed Analysis of the Greedy Algorithm for the Linear Contextual Bandit Problem
https://papers.nips.cc/paper/7491-a-smoothed-analysis-of-the-greedy-algorithm-for-the-linear-contextual-bandit-problem
AUTHORS: Sampath Kannan, Jamie H. Morgenstern, Aaron Roth, Bo Waggoner, Zhiwei Steven Wu
HIGHLIGHT: In this paper, we consider the linear contextual bandit problem and revisit the performance of the greedy algorithm.

206, TITLE: Revisiting Decomposable Submodular Function Minimization with Incidence Relations
https://papers.nips.cc/paper/7492-revisiting-decomposable-submodular-function-minimization-with-incidence-relations
AUTHORS: Pan Li, Olgica Milenkovic
HIGHLIGHT: We introduce a new approach to decomposable submodular function minimization (DSFM) that exploits incidence relations.

207, TITLE: A Practical Algorithm for Distributed Clustering and Outlier Detection
https://papers.nips.cc/paper/7493-a-practical-algorithm-for-distributed-clustering-and-outlier-detection
AUTHORS: Jiecao Chen, Erfan Sadeqi Azer, Qin Zhang
HIGHLIGHT: We propose a simple approach based on constructing small summary for the original dataset.

208, TITLE: Learning to Reconstruct Shapes from Unseen Classes

https://papers.nips.cc/paper/7494-learning-to-reconstruct-shapes-from-unseen-classes
AUTHORS: Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Josh Tenenbaum, Bill Freeman, Jiajun Wu
HIGHLIGHT: Here we present an algorithm, Generalizable Reconstruction (GenRe), designed to capture more generic, class-agnostic shape priors.

209, TITLE: BourGAN: Generative Networks with Metric Embeddings
https://papers.nips.cc/paper/7495-bourgan-generative-networks-with-metric-embeddings
AUTHORS: Chang Xiao, Peilin Zhong, Changxi Zheng
HIGHLIGHT: We use the Gaussian mixture model in tandem with a simple augmentation of the objective function to train GANs.

210, TITLE: Smoothed analysis of the low-rank approach for smooth semidefinite programs
https://papers.nips.cc/paper/7496-smoothed-analysis-of-the-low-rank-approach-for-smooth-semidefinite-programs
AUTHORS: Thomas Pumir, Samy Jelassi, Nicolas Boumal
HIGHLIGHT: To this end, and under similar assumptions, we use smoothed analysis to show that approximate SOSPs for a randomly perturbed objective function are approximate global optima, with k scaling like the square root of the number of constraints (up to log factors).

211, TITLE: Zero-Shot Transfer with Deictic Object-Oriented Representation in Reinforcement Learning
https://papers.nips.cc/paper/7497-zero-shot-transfer-with-deictic-object-oriented-representation-in-reinforcement-learning
AUTHORS: Ofir Marom, Benjamin Rosman
HIGHLIGHT: In this paper we introduce a novel deictic object-oriented framework that has provably efficient learning bounds and can solve a broader range of tasks.

212, TITLE: Overfitting or perfect fitting? Risk bounds for classification and regression rules that interpolate
https://papers.nips.cc/paper/7498-overfitting-or-perfect-fitting-risk-bounds-for-classification-and-regression-rules-that-interpolate
AUTHORS: Mikhail Belkin, Daniel J. Hsu, Partha Mitra
HIGHLIGHT: Many modern machine learning models are trained to achieve zero or near-zero training error in order to obtain near-optimal (but non-zero) test error.

213, TITLE: Breaking the Span Assumption Yields Fast Finite-Sum Minimization
https://papers.nips.cc/paper/7499-breaking-the-span-assumption-yields-fast-finite-sum-minimization
AUTHORS: Robert Hannah, Yanli Liu, Daniel O'Connor, Wotao Yin
HIGHLIGHT: In this paper, we show that SVRG and SARAH can be modified to be fundamentally faster than all of the other standard algorithms that minimize the sum of n smooth functions, such as SAGA, SAG, SDCA, and SDCA without duality.

214, TITLE: Structured Local Minima in Sparse Blind Deconvolution
https://papers.nips.cc/paper/7500-structured-local-minima-in-sparse-blind-deconvolution
AUTHORS: Yuqian Zhang, Han-wen Kuo, John Wright
HIGHLIGHT: We assume the short signal to have unit ℓ^2 norm and cast the blind deconvolution problem as a nonconvex optimization problem over the sphere.

215, TITLE: GIANT: Globally Improved Approximate Newton Method for Distributed Optimization
https://papers.nips.cc/paper/7501-giant-globally-improved-approximate-newton-method-for-distributed-optimization
AUTHORS: Shusen Wang, Farbod Roosta-Khorasani, Peng Xu, Michael W. Mahoney
HIGHLIGHT: For distributed computing environment, we consider the empirical risk minimization problem and propose a distributed and communication-efficient Newton-type optimization method.

216, TITLE: Modelling sparsity, heterogeneity, reciprocity and community structure in temporal interaction data
https://papers.nips.cc/paper/7502-modelling-sparsity-heterogeneity-reciprocity-and-community-structure-in-temporal-interaction-data
AUTHORS: Xenia Miscouridou, Francois Caron, Yee Whye Teh
HIGHLIGHT: We propose a novel class of network models for temporal dyadic interaction data.

217, TITLE: Non-monotone Submodular Maximization in Exponentially Fewer Iterations
https://papers.nips.cc/paper/7503-non-monotone-submodular-maximization-in-exponentially-fewer-iterations
AUTHORS: Eric Balkanski, Adam Breuer, Yaron Singer
HIGHLIGHT: In this paper we consider parallelization for applications whose objective can be expressed as maximizing a non-monotone submodular function under a cardinality constraint.

218, TITLE: MetaGAN: An Adversarial Approach to Few-Shot Learning

https://papers.nips.cc/paper/7504-metagan-an-adversarial-approach-to-few-shot-learning
AUTHORS: Ruixiang ZHANG, Tong Che, Zoubin Ghahramani, Yoshua Bengio, Yangqiu Song
HIGHLIGHT: In this paper, we propose a conceptually simple and general framework called MetaGAN for few-shot learning problems.

219, TITLE: Local Differential Privacy for Evolving Data
https://papers.nips.cc/paper/7505-local-differential-privacy-for-evolving-data
AUTHORS: Matthew Joseph, Aaron Roth, Jonathan Ullman, Bo Waggoner
HIGHLIGHT: In this paper, we introduce a new technique for local differential privacy that makes it possible to maintain up-to-date statistics over time, with privacy guarantees that degrade only in the number of changes in the underlying distribution rather than the number of collection periods.

220, TITLE: Gaussian Process Conditional Density Estimation
https://papers.nips.cc/paper/7506-gaussian-process-conditional-density-estimation
AUTHORS: Vincent Dutordoir, Hugh Salimbeni, James Hensman, Marc Deisenroth
HIGHLIGHT: In this work, we propose to extend the model's input with latent variables and use Gaussian processes (GP) to map this augmented input onto samples from the conditional distribution.

221, TITLE: Meta-Gradient Reinforcement Learning
https://papers.nips.cc/paper/7507-meta-gradient-reinforcement-learning
AUTHORS: Zhongwen Xu, Hado P. van Hasselt, David Silver
HIGHLIGHT: We discuss a gradient-based meta-learning algorithm that is able to adapt the nature of the return, online, whilst interacting and learning from the environment.

222, TITLE: Modular Networks: Learning to Decompose Neural Computation
https://papers.nips.cc/paper/7508-modular-networks-learning-to-decompose-neural-computation
AUTHORS: Louis Kirsch, Julius Kunze, David Barber
HIGHLIGHT: We propose a training algorithm that flexibly chooses neural modules based on the data to be processed.

223, TITLE: Learning to Navigate in Cities Without a Map
https://papers.nips.cc/paper/7509-learning-to-navigate-in-cities-without-a-map
AUTHORS: Piotr Mirowski, Matt Grimes, Mateusz Malinowski, Karl Moritz Hermann, Keith Anderson, Denis Teplyashin, Karen Simonyan, koray kavukcuoglu, Andrew Zisserman, Raia Hadsell
HIGHLIGHT: A key contribution of this paper is an interactive navigation environment that uses Google Street View for its photographic content and worldwide coverage.

224, TITLE: Query Complexity of Bayesian Private Learning
https://papers.nips.cc/paper/7510-query-complexity-of-bayesian-private-learning
AUTHORS: Kuang Xu
HIGHLIGHT: We study the query complexity of Bayesian Private Learning: a learner wishes to locate a random target within an interval by submitting queries, in the presence of an adversary who observes all of her queries but not the responses.

225, TITLE: A theory on the absence of spurious solutions for nonconvex and nonsmooth optimization
https://papers.nips.cc/paper/7511-a-theory-on-the-absence-of-spurious-solutions-for-nonconvex-and-nonsmooth-optimization
AUTHORS: Cedric Josz, Yi Ouyang, Richard Zhang, Javad Lavaei, Somayeh Sojoudi
HIGHLIGHT: We study the set of continuous functions that admit no spurious local optima (i.e. local minima that are not global minima) which we term global functions.

226, TITLE: Recurrent World Models Facilitate Policy Evolution
https://papers.nips.cc/paper/7512-recurrent-world-models-facilitate-policy-evolution
AUTHORS: David Ha, J?rgen Schmidhuber
HIGHLIGHT: Recurrent World Models Facilitate Policy Evolution

227, TITLE: Ridge Regression and Provable Deterministic Ridge Leverage Score Sampling
https://papers.nips.cc/paper/7513-ridge-regression-and-provable-deterministic-ridge-leverage-score-sampling
AUTHORS: Shannon McCurdy
HIGHLIGHT: We provide provable guarantees for deterministic column sampling using ridge leverage scores.

228, TITLE: Wasserstein Variational Inference
https://papers.nips.cc/paper/7514-wasserstein-variational-inference

AUTHORS: Luca Ambrogioni, Umut G??l?, Yagmur G??l?t?rk, Max Hinne, Marcel A. J. van Gerven, Eric Maris
HIGHLIGHT: This paper introduces Wasserstein variational inference, a new form of approximate Bayesian inference based on optimal transport theory.

229, TITLE: How Does Batch Normalization Help Optimization?
https://papers.nips.cc/paper/7515-how-does-batch-normalization-help-optimization
AUTHORS: Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, Aleksander Madry
HIGHLIGHT: In this work, we demonstrate that such distributional stability of layer inputs has little to do with the success of BatchNorm.

230, TITLE: Verifiable Reinforcement Learning via Policy Extraction
https://papers.nips.cc/paper/7516-verifiable-reinforcement-learning-via-policy-extraction
AUTHORS: Osbert Bastani, Yewen Pu, Armando Solar-Lezama
HIGHLIGHT: We propose an approach to verifiable reinforcement learning by training decision tree policies, which can represent complex policies (since they are nonparametric), yet can be efficiently verified using existing techniques (since they are highly structured).

231, TITLE: Leveraged volume sampling for linear regression
https://papers.nips.cc/paper/7517-leveraged-volume-sampling-for-linear-regression
AUTHORS: Michal Derezinski, Manfred K. Warmuth, Daniel J. Hsu
HIGHLIGHT: It is therefore natural to ask if this method offers the optimal unbiased estimate in terms of the number of responses k needed to achieve a 1+epsilon loss approximation.

232, TITLE: Model Agnostic Supervised Local Explanations
https://papers.nips.cc/paper/7518-model-agnostic-supervised-local-explanations
AUTHORS: Gregory Plumb, Denali Molitor, Ameet S. Talwalkar
HIGHLIGHT: We address this challenge in a novel model called MAPLE that uses local linear modeling techniques along with a dual interpretation of random forests (both as a supervised neighborhood approach and as a feature selection method).

233, TITLE: A Linear Speedup Analysis of Distributed Deep Learning with Sparse and Quantized Communication
https://papers.nips.cc/paper/7519-a-linear-speedup-analysis-of-distributed-deep-learning-with-sparse-and-quantized-communication
AUTHORS: Peng Jiang, Gagan Agrawal
HIGHLIGHT: In this paper, we study the convergence rate of distributed SGD for non-convex optimization with two communication reducing strategies: sparse parameter averaging and gradient quantization.

234, TITLE: Active Learning for Non-Parametric Regression Using Purely Random Trees
https://papers.nips.cc/paper/7520-active-learning-for-non-parametric-regression-using-purely-random-trees
AUTHORS: Jack Goetz, Ambuj Tewari, Paul Zimmerman
HIGHLIGHT: In this paper we propose an intuitive tree based active learning algorithm for non-parametric regression with provable improvement over random sampling.

235, TITLE: Tree-to-tree Neural Networks for Program Translation
https://papers.nips.cc/paper/7521-tree-to-tree-neural-networks-for-program-translation
AUTHORS: Xinyun Chen, Chang Liu, Dawn Song
HIGHLIGHT: In this work, we are the first to employ deep neural networks toward tackling this problem.

236, TITLE: Batch-Instance Normalization for Adaptively Style-Invariant Neural Networks
https://papers.nips.cc/paper/7522-batch-instance-normalization-for-adaptively-style-invariant-neural-networks
AUTHORS: Hyeonseob Nam, Hyo-Eun Kim
HIGHLIGHT: Extending this idea to general visual recognition problems, we present Batch-Instance Normalization (BIN) to explicitly normalize unnecessary styles from images.

237, TITLE: Structural Causal Bandits: Where to Intervene?
https://papers.nips.cc/paper/7523-structural-causal-bandits-where-to-intervene
AUTHORS: Sanghack Lee, Elias Bareinboim
HIGHLIGHT: In this paper, we show that whenever the underlying causal model is not taken into account during the decision-making process, the standard strategies of simultaneously intervening on all variables or on all the subsets of the variables may, in general, lead to suboptimal policies, regardless of the number of interventions performed by the agent in the environment.

238, TITLE: Answerer in Questioner's Mind: Information Theoretic Approach to Goal-Oriented Visual Dialog

https://papers.nips.cc/paper/7524-answerer-in-questioners-mind-information-theoretic-approach-to-goal-oriented-visual-dialog
AUTHORS: Sang-Woo Lee, Yu-Jung Heo, Byoung-Tak Zhang
HIGHLIGHT: Motivated by theory of mind, we propose "Answerer in Questioner's Mind" (AQM), a novel information theoretic algorithm for goal-oriented dialog.

239, TITLE: A Unified Feature Disentangler for Multi-Domain Image Translation and Manipulation
https://papers.nips.cc/paper/7525-a-unified-feature-disentangler-for-multi-domain-image-translation-and-manipulation
AUTHORS: Alexander H. Liu, Yen-Cheng Liu, Yu-Ying Yeh, Yu-Chiang Frank Wang
HIGHLIGHT: We present a novel and unified deep learning framework which is capable of learning domain-invariant representation from data across multiple domains.

240, TITLE: Online Learning with an Unknown Fairness Metric
https://papers.nips.cc/paper/7526-online-learning-with-an-unknown-fairness-metric
AUTHORS: Stephen Gillen, Christopher Jung, Michael Kearns, Aaron Roth
HIGHLIGHT: We consider the problem of online learning in the linear contextual bandits setting, but in which there are also strong individual fairness constraints governed by an unknown similarity metric.

241, TITLE: Isolating Sources of Disentanglement in Variational Autoencoders
https://papers.nips.cc/paper/7527-isolating-sources-of-disentanglement-in-variational-autoencoders
AUTHORS: Tian Qi Chen, Xuechen Li, Roger B. Grosse, David K. Duvenaud
HIGHLIGHT: We use this to motivate the beta-TCVAE (Total Correlation Variational Autoencoder) algorithm, a refinement and plug-in replacement of the beta-VAE for learning disentangled representations, requiring no additional hyperparameters during training.

242, TITLE: Contextual bandits with surrogate losses: Margin bounds and efficient algorithms
https://papers.nips.cc/paper/7528-contextual-bandits-with-surrogate-losses-margin-bounds-and-efficient-algorithms
AUTHORS: Dylan J. Foster, Akshay Krishnamurthy
HIGHLIGHT: Using the hinge loss, we derive an efficient algorithm with a $\sqrt{dT}$-type mistake bound against benchmark policies induced by d-dimensional regressors.

243, TITLE: Representation Learning for Treatment Effect Estimation from Observational Data
https://papers.nips.cc/paper/7529-representation-learning-for-treatment-effect-estimation-from-observational-data
AUTHORS: Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, Aidong Zhang
HIGHLIGHT: In this paper, we propose a local similarity preserved individual treatment effect (SITE) estimation method based on deep representation learning.

244, TITLE: Representation Balancing MDPs for Off-policy Policy Evaluation
https://papers.nips.cc/paper/7530-representation-balancing-mdps-for-off-policy-policy-evaluation
AUTHORS: Yao Liu, Omer Gottesman, Aniruddh Raghu, Matthieu Komorowski, Aldo A. Faisal, Finale Doshi-Velez, Emma Brunskill
HIGHLIGHT: We study the problem of off-policy policy evaluation (OPPE) in RL.

245, TITLE: Out of the Box: Reasoning with Graph Convolution Nets for Factual Visual Question Answering
https://papers.nips.cc/paper/7531-out-of-the-box-reasoning-with-graph-convolution-nets-for-factual-visual-question-answering
AUTHORS: Medhini Narasimhan, Svetlana Lazebnik, Alexander Schwing
HIGHLIGHT: We observe that a successive process which considers one fact at a time to form a local decision is sub-optimal.

246, TITLE: Causal Discovery from Discrete Data using Hidden Compact Representation
https://papers.nips.cc/paper/7532-causal-discovery-from-discrete-data-using-hidden-compact-representation
AUTHORS: Ruichu Cai, Jie Qiao, Kun Zhang, Zhenjie Zhang, Zhifeng Hao
HIGHLIGHT: In this paper we make an attempt to find a way to solve this problem by assuming a two-stage causal process: the first stage maps the cause to a hidden variable of a lower cardinality, and the second stage generates the effect from the hidden representation.

247, TITLE: Natasha 2: Faster Non-Convex Optimization Than SGD
https://papers.nips.cc/paper/7533-natasha-2-faster-non-convex-optimization-than-sgd
AUTHORS: Zeyuan Allen-Zhu
HIGHLIGHT: We design a stochastic algorithm to find ε-approximate local minima of any smooth nonconvex function in rate $O(\varepsilon^{-3.25})$, with only oracle access to stochastic gradients.

248, TITLE: Minimax Statistical Learning with Wasserstein distances
https://papers.nips.cc/paper/7534-minimax-statistical-learning-with-wasserstein-distances
AUTHORS: Jaeho Lee, Maxim Raginsky
HIGHLIGHT: In this work, we describe a minimax framework for statistical learning with ambiguity sets given by balls in Wasserstein space.

249, TITLE: Provable Variational Inference for Constrained Log-Submodular Models
https://papers.nips.cc/paper/7535-provable-variational-inference-for-constrained-log-submodular-models
AUTHORS: Josip Djolonga, Stefanie Jegelka, Andreas Krause
HIGHLIGHT: To perform inference in these models we design novel variational inference algorithms, which carefully leverage the combinatorial and probabilistic properties of these objects.

250, TITLE: Learning Hierarchical Semantic Image Manipulation through Structured Representations
https://papers.nips.cc/paper/7536-learning-hierarchical-semantic-image-manipulation-through-structured-representations
AUTHORS: Seunghoon Hong, Xinchen Yan, Thomas S. Huang, Honglak Lee
HIGHLIGHT: In this work, we present a novel hierarchical framework for semantic image manipulation.

251, TITLE: Processing of missing data by neural networks
https://papers.nips.cc/paper/7537-processing-of-missing-data-by-neural-networks
AUTHORS: Marek Smieja, Lukasz Struski, Jacek Tabor, Bartosz Zielinski, Przemyslaw Spurek
HIGHLIGHT: We propose a general, theoretically justified mechanism for processing missing data by neural networks.

252, TITLE: Safe Active Learning for Time-Series Modeling with Gaussian Processes
https://papers.nips.cc/paper/7538-safe-active-learning-for-time-series-modeling-with-gaussian-processes
AUTHORS: Christoph Zimmer, Mona Meister, Duy Nguyen-Tuong
HIGHLIGHT: In this study, we consider the problem of actively learning time-series models while taking given safety constraints into account.

253, TITLE: Optimal Algorithms for Non-Smooth Distributed Optimization in Networks
https://papers.nips.cc/paper/7539-optimal-algorithms-for-non-smooth-distributed-optimization-in-networks
AUTHORS: Kevin Scaman, Francis Bach, Sebastien Bubeck, Laurent Massouli?, Yin Tat Lee
HIGHLIGHT: In this work, we consider the distributed optimization of non-smooth convex functions using a network of computing units.

254, TITLE: Computing Higher Order Derivatives of Matrix and Tensor Expressions
https://papers.nips.cc/paper/7540-computing-higher-order-derivatives-of-matrix-and-tensor-expressions
AUTHORS: Soeren Laue, Matthias Mitterreiter, Joachim Giesen
HIGHLIGHT: Here, we close this fundamental gap and present an algorithmic framework for computing matrix and tensor derivatives that extends seamlessly to higher order derivatives.

255, TITLE: Paraphrasing Complex Network: Network Compression via Factor Transfer
https://papers.nips.cc/paper/7541-paraphrasing-complex-network-network-compression-via-factor-transfer
AUTHORS: Jangho Kim, Seonguk Park, Nojun Kwak
HIGHLIGHT: In this paper, we propose a novel knowledge transfer method which uses convolutional operations to paraphrase teacher's knowledge and to translate it for the student.

256, TITLE: Analytic solution and stationary phase approximation for the Bayesian lasso and elastic net
https://papers.nips.cc/paper/7542-analytic-solution-and-stationary-phase-approximation-for-the-bayesian-lasso-and-elastic-net
AUTHORS: Tom Michoel
HIGHLIGHT: Analytic solution and stationary phase approximation for the Bayesian lasso and elastic net

257, TITLE: Demystifying excessively volatile human learning: A Bayesian persistent prior and a neural approximation
https://papers.nips.cc/paper/7543-demystifying-excessively-volatile-human-learning-a-bayesian-persistent-prior-and-a-neural-approximation
AUTHORS: Chaitanya Ryali, Gautam Reddy, Angela J. Yu
HIGHLIGHT: Here, we examine a neurally plausible algorithm, a special case of leaky integration dynamics we denote as EXP (for exponential filtering), that is significantly simpler than all previously suggested algorithms except for the delta-learning rule, and which far outperforms the delta rule in approximating Bayesian prediction performance.

258, TITLE: Empirical Risk Minimization Under Fairness Constraints

https://papers.nips.cc/paper/7544-empirical-risk-minimization-under-fairness-constraints
AUTHORS: Michele Donini, Luca Oneto, Shai Ben-David, John S. Shawe-Taylor, Massimiliano Pontil
HIGHLIGHT: We present an approach based on empirical risk minimization, which incorporates a fairness constraint into the learning problem.

259, TITLE: Unsupervised Learning of Shape and Pose with Differentiable Point Clouds
https://papers.nips.cc/paper/7545-unsupervised-learning-of-shape-and-pose-with-differentiable-point-clouds
AUTHORS: Eldar Insafutdinov, Alexey Dosovitskiy
HIGHLIGHT: To deal with pose ambiguity, we introduce an ensemble of pose predictors which we then distill to a single "student" model.

260, TITLE: Continuous-time Value Function Approximation in Reproducing Kernel Hilbert Spaces
https://papers.nips.cc/paper/7546-continuous-time-value-function-approximation-in-reproducing-kernel-hilbert-spaces
AUTHORS: Motoya Ohnishi, Masahiro Yukawa, Mikael Johansson, Masashi Sugiyama
HIGHLIGHT: In this paper, we propose a novel framework for model-based continuous-time value function approximation in reproducing kernel Hilbert spaces.

261, TITLE: Gradient Descent Meets Shift-and-Invert Preconditioning for Eigenvector Computation
https://papers.nips.cc/paper/7547-gradient-descent-meets-shift-and-invert-preconditioning-for-eigenvector-computation
AUTHORS: Zhiqiang Xu
HIGHLIGHT: In this work, we adopt an inexact Riemannian gradient descent perspective to investigate this technique on the effect of the step-size scheme.

262, TITLE: Factored Bandits
https://papers.nips.cc/paper/7548-factored-bandits
AUTHORS: Julian Zimmert, Yevgeny Seldin
HIGHLIGHT: We introduce the factored bandits model, which is a framework for learning with limited (bandit) feedback, where actions can be decomposed into a Cartesian product of atomic actions.

263, TITLE: Delta-encoder: an effective sample synthesis method for few-shot object recognition
https://papers.nips.cc/paper/7549-delta-encoder-an-effective-sample-synthesis-method-for-few-shot-object-recognition
AUTHORS: Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, Alex Bronstein
HIGHLIGHT: In this work, we propose a simple yet effective method for few-shot (and one-shot) object recognition.

264, TITLE: Metric on Nonlinear Dynamical Systems with Perron-Frobenius Operators
https://papers.nips.cc/paper/7550-metric-on-nonlinear-dynamical-systems-with-perron-frobenius-operators
AUTHORS: Isao Ishikawa, Keisuke Fujii, Masahiro Ikeda, Yuka Hashimoto, Yoshinobu Kawahara
HIGHLIGHT: In this paper, we develop a general metric for comparing nonlinear dynamical systems that is defined with Perron-Frobenius operators in reproducing kernel Hilbert spaces.

265, TITLE: Learning a High Fidelity Pose Invariant Model for High-resolution Face Frontalization
https://papers.nips.cc/paper/7551-learning-a-high-fidelity-pose-invariant-model-for-high-resolution-face-frontalization
AUTHORS: Jie Cao, Yibo Hu, Hongwen Zhang, Ran He, Zhenan Sun
HIGHLIGHT: This paper proposes a High Fidelity Pose Invariant Model (HF-PIM) to produce photographic and identity-preserving results.

266, TITLE: Mirrored Langevin Dynamics
https://papers.nips.cc/paper/7552-mirrored-langevin-dynamics
AUTHORS: Ya-Ping Hsieh, Ali Kavis, Paul Rolland, Volkan Cevher
HIGHLIGHT: We propose a unified framework, which is inspired by the classical mirror descent, to derive novel first-order sampling schemes.

267, TITLE: Moonshine: Distilling with Cheap Convolutions
https://papers.nips.cc/paper/7553-moonshine-distilling-with-cheap-convolutions
AUTHORS: Elliot J. Crowley, Gavin Gray, Amos J. Storkey
HIGHLIGHT: We propose structural model distillation for memory reduction using a strategy that produces a student architecture that is a simple transformation of the teacher architecture: no redesign is needed, and the same hyperparameters can be used.

268, TITLE: Stochastic Cubic Regularization for Fast Nonconvex Optimization
https://papers.nips.cc/paper/7554-stochastic-cubic-regularization-for-fast-nonconvex-optimization
AUTHORS: Nilesh Tripuraneni, Mitchell Stern, Chi Jin, Jeffrey Regier, Michael I. Jordan
HIGHLIGHT: This paper proposes a stochastic variant of a classic algorithm---the cubic-regularized Newton method [Nesterov and Polyak].

269, TITLE: Adaptation to Easy Data in Prediction with Limited Advice
https://papers.nips.cc/paper/7555-adaptation-to-easy-data-in-prediction-with-limited-advice
AUTHORS: Tobias Thune, Yevgeny Seldin
HIGHLIGHT: We derive an online learning algorithm with improved regret guarantees for ``easy'' loss sequences.

270, TITLE: Differentially Private Bayesian Inference for Exponential Families
https://papers.nips.cc/paper/7556-differentially-private-bayesian-inference-for-exponential-families
AUTHORS: Garrett Bernstein, Daniel R. Sheldon
HIGHLIGHT: We present the first method for private Bayesian inference in exponential families that properly accounts for noise introduced by the privacy mechanism.

271, TITLE: Playing hard exploration games by watching YouTube
https://papers.nips.cc/paper/7557-playing-hard-exploration-games-by-watching-youtube
AUTHORS: Yusuf Aytar, Tobias Pfaff, David Budden, Thomas Paine, Ziyu Wang, Nando de Freitas
HIGHLIGHT: Here we propose a method that overcomes these limitations in two stages.

272, TITLE: Dialog-to-Action: Conversational Question Answering Over a Large-Scale Knowledge Base
https://papers.nips.cc/paper/7558-dialog-to-action-conversational-question-answering-over-a-large-scale-knowledge-base
AUTHORS: Daya Guo, Duyu Tang, Nan Duan, Ming Zhou, Jian Yin
HIGHLIGHT: We present an approach to map utterances in conversation to logical forms, which will be executed on a large-scale knowledge base.

273, TITLE: Norm-Ranging LSH for Maximum Inner Product Search
https://papers.nips.cc/paper/7559-norm-ranging-lsh-for-maximum-inner-product-search
AUTHORS: Xiao Yan, Jinfeng Li, Xinyan Dai, Hongzhi Chen, James Cheng
HIGHLIGHT: We propose NORM-RANGING LSH, which addresses the excessive normalization problem caused by long tails by partitioning a dataset into sub-datasets and building a hash index for each sub-dataset independently.

274, TITLE: Optimization over Continuous and Multi-dimensional Decisions with Observational Data
https://papers.nips.cc/paper/7560-optimization-over-continuous-and-multi-dimensional-decisions-with-observational-data
AUTHORS: Dimitris Bertsimas, Christopher McCord
HIGHLIGHT: We propose a novel algorithmic framework that is tractable, asymptotically consistent, and superior to comparable methods on example problems.

275, TITLE: Fast Estimation of Causal Interactions using Wold Processes
https://papers.nips.cc/paper/7561-fast-estimation-of-causal-interactions-using-wold-processes
AUTHORS: Flavio Figueiredo, Guilherme Resende Borges, Pedro O.S. Vaz de Melo, Renato Assun??o
HIGHLIGHT: We here focus on the task of learning Granger causality matrices for multivariate point processes.

276, TITLE: When do random forests fail?
https://papers.nips.cc/paper/7562-when-do-random-forests-fail
AUTHORS: Cheng Tang, Damien Garreau, Ulrike von Luxburg
HIGHLIGHT: In this paper, we consider various tree constructions and examine how the choice of parameters affects the generalization error of the resulting random forests as the sample size goes to infinity.

277, TITLE: Near Optimal Exploration-Exploitation in Non-Communicating Markov Decision Processes
https://papers.nips.cc/paper/7563-near-optimal-exploration-exploitation-in-non-communicating-markov-decision-processes
AUTHORS: Ronan Fruit, Matteo Pirotta, Alessandro Lazaric
HIGHLIGHT: In this paper, we introduce TUCRL, the first algorithm able to perform efficient exploration-exploitation in any finite Markov Decision Process (MDP) without requiring any form of prior knowledge.

278, TITLE: Optimistic optimization of a Brownian
https://papers.nips.cc/paper/7564-optimistic-optimization-of-a-brownian
AUTHORS: Jean-Bastien Grill, Michal Valko, Remi Munos

HIGHLIGHT: Given W, our goal is to return an ϵ-approximation of its maximum using the smallest possible number of function evaluations, the sample complexity of the algorithm.

279, TITLE: Practical Methods for Graph Two-Sample Testing
https://papers.nips.cc/paper/7565-practical-methods-for-graph-two-sample-testing
AUTHORS: Debarghya Ghoshdastidar, Ulrike von Luxburg
HIGHLIGHT: In this paper, we consider the problem of two-sample testing of large graphs.

280, TITLE: NAIS-Net: Stable Deep Networks from Non-Autonomous Differential Equations
https://papers.nips.cc/paper/7566-nais-net-stable-deep-networks-from-non-autonomous-differential-equations
AUTHORS: Marco Ciccone, Marco Gallieri, Jonathan Masci, Christian Osendorfer, Faustino Gomez
HIGHLIGHT: This paper introduces Non-Autonomous Input-Output Stable Network (NAIS-Net), a very deep architecture where each stacked processing block is derived from a time-invariant non-autonomous dynamical system.

281, TITLE: On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport
https://papers.nips.cc/paper/7567-on-the-global-convergence-of-gradient-descent-for-over-parameterized-models-using-optimal-transport
AUTHORS: L?na?c Chizat, Francis Bach
HIGHLIGHT: For these problems, we study a simple minimization method: the unknown measure is discretized into a mixture of particles and a continuous-time gradient descent is performed on their weights and positions.

282, TITLE: Constructing Deep Neural Networks by Bayesian Network Structure Learning
https://papers.nips.cc/paper/7568-constructing-deep-neural-networks-by-bayesian-network-structure-learning
AUTHORS: Raanan Y. Rohekar, Shami Nisimov, Yaniv Gurwicz, Guy Koren, Gal Novik
HIGHLIGHT: We introduce a principled approach for unsupervised structure learning of deep neural networks.

283, TITLE: Weakly Supervised Dense Event Captioning in Videos
https://papers.nips.cc/paper/7569-weakly-supervised-dense-event-captioning-in-videos
AUTHORS: Xuguang Duan, Wenbing Huang, Chuang Gan, Jingdong Wang, Wenwu Zhu, Junzhou Huang
HIGHLIGHT: This paper formulates a new problem: weakly supervised dense event captioning, which does not require temporal segment annotations for model training.

284, TITLE: Faithful Inversion of Generative Models for Effective Amortized Inference
https://papers.nips.cc/paper/7570-faithful-inversion-of-generative-models-for-effective-amortized-inference
AUTHORS: Stefan Webb, Adam Golinski, Rob Zinkov, Siddharth Narayanaswamy, Tom Rainforth, Yee Whye Teh, Frank Wood
HIGHLIGHT: We introduce an algorithm for faithfully, and minimally, inverting the graphical model structure of any generative model.

285, TITLE: From Stochastic Planning to Marginal MAP
https://papers.nips.cc/paper/7571-from-stochastic-planning-to-marginal-map
AUTHORS: Hao Cui, Radu Marinescu, Roni Khardon
HIGHLIGHT: We introduce a new reduction from MMAP to maximum expected utility problems which are suitable for the symbolic computation in SOGBOFA.

286, TITLE: On Binary Classification in Extreme Regions
https://papers.nips.cc/paper/7572-on-binary-classification-in-extreme-regions
AUTHORS: Hamid JALALZAI, Stephan Cl?men?on, Anne Sabourin
HIGHLIGHT: It is the purpose of this paper to develop a general framework for classification in the extremes.

287, TITLE: Near-Optimal Policies for Dynamic Multinomial Logit Assortment Selection Models
https://papers.nips.cc/paper/7573-near-optimal-policies-for-dynamic-multinomial-logit-assortment-selection-models
AUTHORS: Yining Wang, Xi Chen, Yuan Zhou
HIGHLIGHT: In this paper we consider the dynamic assortment selection problem under an uncapacitated multinomial-logit (MNL) model.

288, TITLE: Q-learning with Nearest Neighbors
https://papers.nips.cc/paper/7574-q-learning-with-nearest-neighbors
AUTHORS: Devavrat Shah, Qiaomin Xie

HIGHLIGHT: We consider the Nearest Neighbor Q-Learning (NNQL) algorithm to learn the optimal Q function using nearest neighbor regression method.

289, TITLE: Global Convergence of Langevin Dynamics Based Algorithms for Nonconvex Optimization
https://papers.nips.cc/paper/7575-global-convergence-of-langevin-dynamics-based-algorithms-for-nonconvex-optimization
AUTHORS: Pan Xu, Jinghui Chen, Difan Zou, Quanquan Gu
HIGHLIGHT: We present a unified framework to analyze the global convergence of Langevin dynamics based algorithms for nonconvex finite-sum optimization with n component functions.

290, TITLE: Asymptotic optimality of adaptive importance sampling
https://papers.nips.cc/paper/7576-asymptotic-optimality-of-adaptive-importance-sampling
AUTHORS: Fran?ois Portier, Bernard Delyon
HIGHLIGHT: The very fundamental question raised in this paper concerns the behavior of empirical sums based on AIS.

291, TITLE: Learning latent variable structured prediction models with Gaussian perturbations
https://papers.nips.cc/paper/7577-learning-latent-variable-structured-prediction-models-with-gaussian-perturbations
AUTHORS: Kevin Bello, Jean Honorio
HIGHLIGHT: Recent work has proposed the use of the maximum loss over random structured outputs sampled independently from some proposal distribution, with theoretical guarantees.

292, TITLE: The Nearest Neighbor Information Estimator is Adaptively Near Minimax Rate-Optimal
https://papers.nips.cc/paper/7578-the-nearest-neighbor-information-estimator-is-adaptively-near-minimax-rate-optimal
AUTHORS: Jiantao Jiao, Weihao Gao, Yanjun Han
HIGHLIGHT: We analyze the Kozachenko–Leonenko (KL) fixed k-nearest neighbor estimator for the differential entropy.

293, TITLE: Deep Reinforcement Learning of Marked Temporal Point Processes
https://papers.nips.cc/paper/7579-deep-reinforcement-learning-of-marked-temporal-point-processes
AUTHORS: Utkarsh Upadhyay, Abir De, Manuel Gomez Rodriguez
HIGHLIGHT: In this paper, we address the above problem from the perspective of deep reinforcement learning of marked temporal point processes, where both the actions taken by an agent and the feedback it receives from the environment are asynchronous stochastic discrete events characterized using marked temporal point processes.

294, TITLE: Evidential Deep Learning to Quantify Classification Uncertainty
https://papers.nips.cc/paper/7580-evidential-deep-learning-to-quantify-classification-uncertainty
AUTHORS: Murat Sensoy, Lance Kaplan, Melih Kandemir
HIGHLIGHT: Orthogonally to Bayesian neural nets that indirectly infer prediction uncertainty through weight uncertainties, we propose explicit modeling of the same using the theory of subjective logic.

295, TITLE: Parsimonious Bayesian deep networks
https://papers.nips.cc/paper/7581-parsimonious-bayesian-deep-networks
AUTHORS: Mingyuan Zhou
HIGHLIGHT: Combining Bayesian nonparametrics and a forward model selection strategy, we construct parsimonious Bayesian deep networks (PBDNs) that infer capacity-regularized network architectures from the data and require neither cross-validation nor fine-tuning when training the model.

296, TITLE: Single-Agent Policy Tree Search With Guarantees
https://papers.nips.cc/paper/7582-single-agent-policy-tree-search-with-guarantees
AUTHORS: Laurent Orseau, Levi Lelis, Tor Lattimore, Theophane Weber
HIGHLIGHT: We introduce two novel tree search algorithms that use a policy to guide search.

297, TITLE: Semi-crowdsourced Clustering with Deep Generative Models
https://papers.nips.cc/paper/7583-semi-crowdsourced-clustering-with-deep-generative-models
AUTHORS: Yucen Luo, TIAN TIAN, Jiaxin Shi, Jun Zhu, Bo Zhang
HIGHLIGHT: We propose a new approach that includes a deep generative model (DGM) to characterize low-level features of the data, and a statistical relational model for noisy pairwise annotations on its subset.

298, TITLE: The committee machine: Computational to statistical gaps in learning a two-layers neural network
https://papers.nips.cc/paper/7584-the-committee-machine-computational-to-statistical-gaps-in-learning-a-two-layers-neural-network
AUTHORS: Benjamin Aubin, Antoine Maillard, jean barbier, Florent Krzakala, Nicolas Macris, Lenka Zdeborov?

HIGHLIGHT: In this contribution, we provide a rigorous justification of these approaches for a two-layers neural network model called the committee machine.

299, TITLE: Realistic Evaluation of Deep Semi-Supervised Learning Algorithms
https://papers.nips.cc/paper/7585-realistic-evaluation-of-deep-semi-supervised-learning-algorithms
AUTHORS: Avital Oliver, Augustus Odena, Colin A. Raffel, Ekin Dogus Cubuk, Ian Goodfellow
HIGHLIGHT: Semi-supervised learning (SSL) provides a powerful framework for leveraging unlabeled data when labels are limited or expensive to obtain.

300, TITLE: Contextual Combinatorial Multi-armed Bandits with Volatile Arms and Submodular Reward
https://papers.nips.cc/paper/7586-contextual-combinatorial-multi-armed-bandits-with-volatile-arms-and-submodular-reward
AUTHORS: Lixing Chen, Jie Xu, Zhuo Lu
HIGHLIGHT: In this paper, we study the stochastic contextual combinatorial multi-armed bandit (CC-MAB) framework that is tailored for volatile arms and submodular reward functions.

301, TITLE: Training deep learning based denoisers without ground truth data
https://papers.nips.cc/paper/7587-training-deep-learning-based-denoisers-without-ground-truth-data
AUTHORS: Shakarim Soltanayev, Se Young Chun
HIGHLIGHT: In this article, we propose a method based on Stein's unbiased risk estimator (SURE) for training deep neural network denoisers only based on the use of noisy images.

302, TITLE: Re-evaluating evaluation
https://papers.nips.cc/paper/7588-re-evaluating-evaluation
AUTHORS: David Balduzzi, Karl Tuyls, Julien Perolat, Thore Graepel
HIGHLIGHT: In this paper we take a step back and propose Nash averaging.

303, TITLE: Deep, complex, invertible networks for inversion of transmission effects in multimode optical fibres
https://papers.nips.cc/paper/7589-deep-complex-invertible-networks-for-inversion-of-transmission-effects-in-multimode-optical-fibres
AUTHORS: Ois?n Moran, Piergiorgio Caramazza, Daniele Faccio, Roderick Murray-Smith
HIGHLIGHT: A {\it unitary regularisation} approach for complex-weighted networks is proposed which performs well in robustly inverting the fibre transmission matrix, which fits well with the physical theory.
We generated experimental data based on collections of optical fibre responses to greyscale input images generated with coherent light, by measuring only image amplitude (not amplitude and phase as is typical) at the output of \SI{1}{\metre} and \SI{10}{\metre} long, \SI{105}{\micro\metre} diameter multimode fibre.

304, TITLE: Multivariate Convolutional Sparse Coding for Electromagnetic Brain Signals
https://papers.nips.cc/paper/7590-multivariate-convolutional-sparse-coding-for-electromagnetic-brain-signals
AUTHORS: Tom Dupr? la Tour, Thomas Moreau, Mainak Jas, Alexandre Gramfort
HIGHLIGHT: In this paper, we propose to learn dedicated representations of such recordings using a multivariate convolutional sparse coding (CSC) algorithm.

305, TITLE: Data-Efficient Hierarchical Reinforcement Learning
https://papers.nips.cc/paper/7591-data-efficient-hierarchical-reinforcement-learning
AUTHORS: Ofir Nachum, Shixiang (Shane) Gu, Honglak Lee, Sergey Levine
HIGHLIGHT: In this paper, we study how we can develop HRL algorithms that are general, in that they do not make onerous additional assumptions beyond standard RL algorithms, and efficient, in the sense that they can be used with modest numbers of interaction samples, making them suitable for real-world problems such as robotic control.

306, TITLE: Speaker-Follower Models for Vision-and-Language Navigation
https://papers.nips.cc/paper/7592-speaker-follower-models-for-vision-and-language-navigation
AUTHORS: Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, Trevor Darrell
HIGHLIGHT: Here we describe an approach to vision-and-language navigation that addresses both these issues with an embedded speaker model.

307, TITLE: Inequity aversion improves cooperation in intertemporal social dilemmas
https://papers.nips.cc/paper/7593-inequity-aversion-improves-cooperation-in-intertemporal-social-dilemmas
AUTHORS: Edward Hughes, Joel Z. Leibo, Matthew Phillips, Karl Tuyls, Edgar Due?ez-Guzman, Antonio Garc?a Casta?eda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, Heather Roff, Thore Graepel
HIGHLIGHT: Here we extend this idea to Markov games and show that it promotes cooperation in several types of sequential social dilemma, via a profitable interaction with policy learnability.

308, TITLE: Learning Gaussian Processes by Minimizing PAC-Bayesian Generalization Bounds
https://papers.nips.cc/paper/7594-learning-gaussian-processes-by-minimizing-pac-bayesian-generalization-bounds
AUTHORS: David Reeb, Andreas Doerr, Sebastian Gerwinn, Barbara Rakitsch
HIGHLIGHT: To this end, we propose a method to learn GPs and their sparse approximations by directly optimizing a PAC-Bayesian bound on their generalization performance, instead of maximizing the marginal likelihood.

309, TITLE: Probabilistic Matrix Factorization for Automated Machine Learning
https://papers.nips.cc/paper/7595-probabilistic-matrix-factorization-for-automated-machine-learning
AUTHORS: Nicolo Fusi, Rishit Sheth, Melih Elibol
HIGHLIGHT: In this paper, we propose to solve this meta-learning task by combining ideas from collaborative filtering and Bayesian optimization.

310, TITLE: Stochastic Spectral and Conjugate Descent Methods
https://papers.nips.cc/paper/7596-stochastic-spectral-and-conjugate-descent-methods
AUTHORS: Dmitry Kovalev, Peter Richtarik, Eduard Gorbunov, Elnur Gasanov
HIGHLIGHT: In this paper we introduce a fundamentally new type of acceleration strategy for RCD based on the augmentation of the set of coordinate directions by a few spectral or conjugate directions.

311, TITLE: Recurrent Relational Networks
https://papers.nips.cc/paper/7597-recurrent-relational-networks
AUTHORS: Rasmus Palm, Ulrich Paquet, Ole Winther
HIGHLIGHT: We introduce the recurrent relational net- work, a general purpose module that operates on a graph representation of objects.

312, TITLE: But How Does It Work in Theory? Linear SVM with Random Features
https://papers.nips.cc/paper/7598-but-how-does-it-work-in-theory-linear-svm-with-random-features
AUTHORS: Yitong Sun, Anna Gilbert, Ambuj Tewari
HIGHLIGHT: We prove that, under low noise assumptions, the support vector machine with $N\ll m$ random features (RFSVM) can achieve the learning rate faster than $O(1/\sqrt{m})$ on a training set with m samples when an optimized feature map is used.

313, TITLE: Learning to Optimize Tensor Programs
https://papers.nips.cc/paper/7599-learning-to-optimize-tensor-programs
AUTHORS: Tianqi Chen, Lianmin Zheng, Eddie Yan, Ziheng Jiang, Thierry Moreau, Luis Ceze, Carlos Guestrin, Arvind Krishnamurthy
HIGHLIGHT: We introduce a learning-based framework to optimize tensor programs for deep learning workloads.

314, TITLE: Boosting Black Box Variational Inference
https://papers.nips.cc/paper/7600-boosting-black-box-variational-inference
AUTHORS: Francesco Locatello, Gideon Dresdner, Rajiv Khanna, Isabel Valera, Gunnar Raetsch
HIGHLIGHT: On the theoretical side, we show that boosting VI satisfies a relaxed smoothness assumption which is sufficient for the convergence of the functional Frank-Wolfe (FW) algorithm.

315, TITLE: Nearly tight sample complexity bounds for learning mixtures of Gaussians via sample compression schemes
https://papers.nips.cc/paper/7601-nearly-tight-sample-complexity-bounds-for-learning-mixtures-of-gaussians-via-sample-compression-schemes
AUTHORS: Hassan Ashtiani, Shai Ben-David, Nicholas Harvey, Christopher Liaw, Abbas Mehrabian, Yaniv Plan
HIGHLIGHT: We prove that ?(k d^2 / ?^2) samples are necessary and sufficient for learning a mixture of k Gaussians in R^d, up to error ? in total variation distance.

316, TITLE: Actor-Critic Policy Optimization in Partially Observable Multiagent Environments
https://papers.nips.cc/paper/7602-actor-critic-policy-optimization-in-partially-observable-multiagent-environments
AUTHORS: Sriram Srinivasan, Marc Lanctot, Vinicius Zambaldi, Julien Perolat, Karl Tuyls, Remi Munos, Michael Bowling
HIGHLIGHT: In this paper, we examine the role of these policy gradient and actor-critic algorithms in partially-observable multiagent environments.

317, TITLE: Step Size Matters in Deep Learning
https://papers.nips.cc/paper/7603-step-size-matters-in-deep-learning

AUTHORS: Kamil Nar, Shankar Sastry
HIGHLIGHT: To elucidate the effects of the step size on training of neural networks, we study the gradient descent algorithm as a discrete-time dynamical system, and by analyzing the Lyapunov stability of different solutions, we show the relationship between the step size of the algorithm and the solutions that can be obtained with this algorithm.

318, TITLE: Derivative Estimation in Random Design
https://papers.nips.cc/paper/7604-derivative-estimation-in-random-design
AUTHORS: Yu Liu, Kris De Brabanter
HIGHLIGHT: We propose a nonparametric derivative estimation method for random design without having to estimate the regression function.

319, TITLE: Zeroth-order (Non)-Convex Stochastic Optimization via Conditional Gradient and Gradient Updates
https://papers.nips.cc/paper/7605-zeroth-order-non-convex-stochastic-optimization-via-conditional-gradient-and-gradient-updates
AUTHORS: Krishnakumar Balasubramanian, Saeed Ghadimi
HIGHLIGHT: In this paper, we propose and analyze zeroth-order stochastic approximation algorithms for nonconvex and convex optimization.

320, TITLE: Latent Gaussian Activity Propagation: Using Smoothness and Structure to Separate and Localize Sounds in Large Noisy Environments
https://papers.nips.cc/paper/7606-latent-gaussian-activity-propagation-using-smoothness-and-structure-to-separate-and-localize-sounds-in-large-noisy-environments
AUTHORS: Daniel Johnson, Daniel Gorelik, Ross E. Mawhorter, Kyle Suver, Weiqing Gu, Steven Xing, Cody Gabriel, Peter Sankhagowit
HIGHLIGHT: We present an approach for simultaneously separating and localizing multiple sound sources using recorded microphone data.

321, TITLE: Hybrid-MST: A Hybrid Active Sampling Strategy for Pairwise Preference Aggregation
https://papers.nips.cc/paper/7607-hybrid-mst-a-hybrid-active-sampling-strategy-for-pairwise-preference-aggregation
AUTHORS: JING LI, Rafal Mantiuk, Junle Wang, Suiyi Ling, Patrick Le Callet
HIGHLIGHT: In this paper we present a hybrid active sampling strategy for pairwise preference aggregation, which aims at recovering the underlying rating of the test candidates from sparse and noisy pairwise labeling.

322, TITLE: Infinite-Horizon Gaussian Processes
https://papers.nips.cc/paper/7608-infinite-horizon-gaussian-processes
AUTHORS: Arno Solin, James Hensman, Richard E. Turner
HIGHLIGHT: Gaussian processes provide a flexible framework for forecasting, removing noise, and interpreting long temporal datasets.

323, TITLE: Dimensionality Reduction for Stationary Time Series via Stochastic Nonconvex Optimization
https://papers.nips.cc/paper/7609-dimensionality-reduction-for-stationary-time-series-via-stochastic-nonconvex-optimization
AUTHORS: Minshuo Chen, Lin Yang, Mengdi Wang, Tuo Zhao
HIGHLIGHT: Specifically, our goal is to estimate the principle component of time series data with respect to the covariance matrix of the stationary distribution.

324, TITLE: Sequence-to-Segment Networks for Segment Detection
https://papers.nips.cc/paper/7610-sequence-to-segment-networks-for-segment-detection
AUTHORS: Zijun Wei, Boyu Wang, Minh Hoai Nguyen, Jianming Zhang, Zhe Lin, Xiaohui Shen, Radomir Mech, Dimitris Samaras
HIGHLIGHT: To address this problem, we propose the Sequence-to-Segment Network (S^{2}N), a novel end-to-end sequential encoder-decoder architecture.

325, TITLE: Scaling the Poisson GLM to massive neural datasets through polynomial approximations
https://papers.nips.cc/paper/7611-scaling-the-poisson-glm-to-massive-neural-datasets-through-polynomial-approximations
AUTHORS: David Zoltowski, Jonathan W. Pillow
HIGHLIGHT: We introduce an adaptive procedure to select the polynomial approximation interval and show that the resulting method allows for efficient and accurate inference and regularization of high-dimensional parameters.

326, TITLE: Multiplicative Weights Updates with Constant Step-Size in Graphical Constant-Sum Games
https://papers.nips.cc/paper/7612-multiplicative-weights-updates-with-constant-step-size-in-graphical-constant-sum-games
AUTHORS: Yun Kuen Cheung

HIGHLIGHT: We show that this is false in the context of graphical constant-sum games, which include two-person zero-sum games as special cases.

327, TITLE: Why Is My Classifier Discriminatory?
https://papers.nips.cc/paper/7613-why-is-my-classifier-discriminatory
AUTHORS: Irene Chen, Fredrik D. Johansson, David Sontag
HIGHLIGHT: In this work, we argue that the fairness of predictions should be evaluated in context of the data, and that unfairness induced by inadequate samples sizes or unmeasured predictive variables should be addressed through data collection, rather than by constraining the model.

328, TITLE: Multi-Layered Gradient Boosting Decision Trees
https://papers.nips.cc/paper/7614-multi-layered-gradient-boosting-decision-trees
AUTHORS: Ji Feng, Yang Yu, Zhi-Hua Zhou
HIGHLIGHT: In this work, we propose the multi-layered GBDT forest (mGBDTs), with an explicit emphasis on exploring the ability to learn hierarchical distributed representations by stacking several layers of regression GBDTs as its building block.

329, TITLE: Learn What Not to Learn: Action Elimination with Deep Reinforcement Learning
https://papers.nips.cc/paper/7615-learn-what-not-to-learn-action-elimination-with-deep-reinforcement-learning
AUTHORS: Tom Zahavy, Matan Haroush, Nadav Merlis, Daniel J. Mankowitz, Shie Mannor
HIGHLIGHT: In this work, we propose the Action-Elimination Deep Q-Network (AE-DQN) architecture that combines a Deep RL algorithm with an Action Elimination Network (AEN) that eliminates sub-optimal actions.

330, TITLE: Communication Efficient Parallel Algorithms for Optimization on Manifolds
https://papers.nips.cc/paper/7616-communication-efficient-parallel-algorithms-for-optimization-on-manifolds
AUTHORS: Bayan Saparbayeva, Michael Zhang, Lizhen Lin
HIGHLIGHT: Our work aims to fill a critical gap in the literature by generalizing parallel inference algorithms to optimization on manifolds.

331, TITLE: Neural Code Comprehension: A Learnable Representation of Code Semantics
https://papers.nips.cc/paper/7617-neural-code-comprehension-a-learnable-representation-of-code-semantics
AUTHORS: Tal Ben-Nun, Alice Shoshana Jakobovits, Torsten Hoefler
HIGHLIGHT: In this paper, we propose a novel processing technique to learn code semantics, and apply it to a variety of program analysis tasks.

332, TITLE: Tight Bounds for Collaborative PAC Learning via Multiplicative Weights
https://papers.nips.cc/paper/7618-tight-bounds-for-collaborative-pac-learning-via-multiplicative-weights
AUTHORS: Jiecao Chen, Qin Zhang, Yuan Zhou
HIGHLIGHT: We study the collaborative PAC learning problem recently proposed in Blum et al.~\cite{BHPQ17}, in which we have k players and they want to learn a target function collaboratively, such that the learned function approximates the target function well on all players' distributions simultaneously.

333, TITLE: BinGAN: Learning Compact Binary Descriptors with a Regularized GAN
https://papers.nips.cc/paper/7619-bingan-learning-compact-binary-descriptors-with-a-regularized-gan
AUTHORS: Maciej Zieba, Piotr Semberecki, Tarek El-Gaaly, Tomasz Trzcinski
HIGHLIGHT: In this paper, we propose a novel regularization method for Generative Adversarial Networks that allows the model to learn discriminative yet compact binary representations of image patches (image descriptors).

334, TITLE: Modern Neural Networks Generalize on Small Data Sets
https://papers.nips.cc/paper/7620-modern-neural-networks-generalize-on-small-data-sets
AUTHORS: Matthew Olson, Abraham Wyner, Richard Berk
HIGHLIGHT: In this paper, we use a linear program to empirically decompose fitted neural networks into ensembles of low-bias sub-networks.

335, TITLE: Escaping Saddle Points in Constrained Optimization
https://papers.nips.cc/paper/7621-escaping-saddle-points-in-constrained-optimization
AUTHORS: Aryan Mokhtari, Asuman Ozdaglar, Ali Jadbabaie
HIGHLIGHT: In this paper, we study the problem of escaping from saddle points in smooth nonconvex optimization problems subject to a convex set $\mathcal{C}$.

336, TITLE: Adversarial Attacks on Stochastic Bandits

https://papers.nips.cc/paper/7622-adversarial-attacks-on-stochastic-bandits
AUTHORS: Kwang-Sung Jun, Lihong Li, Yuzhe Ma, Jerry Zhu
HIGHLIGHT: We propose the first attack against two popular bandit algorithms: ϵ-greedy and UCB, \emph{without} knowledge of the mean rewards.

337, TITLE: Optimal Subsampling with Influence Functions
https://papers.nips.cc/paper/7623-optimal-subsampling-with-influence-functions
AUTHORS: Daniel Ting, Eric Brochu
HIGHLIGHT: We show that the concept of an asymptotically linear estimator and the associated influence function leads to asymptotically optimal sampling probabilities for a wide class of popular models.

338, TITLE: A Bandit Approach to Sequential Experimental Design with False Discovery Control
https://papers.nips.cc/paper/7624-a-bandit-approach-to-sequential-experimental-design-with-false-discovery-control
AUTHORS: Kevin G. Jamieson, Lalit Jain
HIGHLIGHT: We propose a new adaptive sampling approach to multiple testing which aims to maximize statistical power while ensuring anytime false discovery control.

339, TITLE: Equality of Opportunity in Classification: A Causal Approach
https://papers.nips.cc/paper/7625-equality-of-opportunity-in-classification-a-causal-approach
AUTHORS: Junzhe Zhang, Elias Bareinboim
HIGHLIGHT: The goal of this paper is to develop a principled approach to connect the statistical disparities characterized by the EO and the underlying, elusive, and frequently unobserved, causal mechanisms that generated such inequality.

340, TITLE: Towards Understanding Acceleration Tradeoff between Momentum and Asynchrony in Nonconvex Stochastic Optimization
https://papers.nips.cc/paper/7626-towards-understanding-acceleration-tradeoff-between-momentum-and-asynchrony-in-nonconvex-stochastic-optimization
AUTHORS: Tianyi Liu, Shiyang Li, Jianping Shi, Enlu Zhou, Tuo Zhao
HIGHLIGHT: Therefore, we propose to analyze the algorithm through a simpler but nontrivial nonconvex problems --- streaming PCA.

341, TITLE: Unsupervised Attention-guided Image-to-Image Translation
https://papers.nips.cc/paper/7627-unsupervised-attention-guided-image-to-image-translation
AUTHORS: Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, Kwang In Kim
HIGHLIGHT: Motivated by the important role of attention in human perception, we tackle this limitation by introducing unsupervised attention mechanisms which are jointly adversarially trained with the generators and discriminators.

342, TITLE: Inferring Networks From Random Walk-Based Node Similarities
https://papers.nips.cc/paper/7628-inferring-networks-from-random-walk-based-node-similarities
AUTHORS: Jeremy Hoskins, Cameron Musco, Christopher Musco, Babis Tsourakakis
HIGHLIGHT: In this work we consider a privacy threat to a social network in which an attacker has access to a subset of random walk-based node similarities, such as effective resistances (i.e., commute times) or personalized PageRank scores.

343, TITLE: NEON2: Finding Local Minima via First-Order Oracles
https://papers.nips.cc/paper/7629-neon2-finding-local-minima-via-first-order-oracles
AUTHORS: Zeyuan Allen-Zhu, Yuanzhi Li
HIGHLIGHT: We propose a reduction for non-convex optimization that can (1) turn an stationary-point finding algorithm into an local-minimum finding one, and (2) replace the Hessian-vector product computations with only gradient computations.

344, TITLE: Zeroth-Order Stochastic Variance Reduction for Nonconvex Optimization
https://papers.nips.cc/paper/7630-zeroth-order-stochastic-variance-reduction-for-nonconvex-optimization
AUTHORS: Sijia Liu, Bhavya Kailkhura, Pin-Yu Chen, Paishun Ting, Shiyu Chang, Lisa Amini
HIGHLIGHT: To mitigate this error, we propose two accelerated versions of ZO-SVRG utilizing variance reduced gradient estimators, which achieve the best rate known for ZO stochastic optimization (in terms of iterations).

345, TITLE: Online Structured Laplace Approximations for Overcoming Catastrophic Forgetting
https://papers.nips.cc/paper/7631-online-structured-laplace-approximations-for-overcoming-catastrophic-forgetting
AUTHORS: Hippolyt Ritter, Aleksandar Botev, David Barber
HIGHLIGHT: We introduce the Kronecker factored online Laplace approximation for overcoming catastrophic forgetting in neural networks.

346, TITLE: DeepProbLog: Neural Probabilistic Logic Programming
https://papers.nips.cc/paper/7632-deepproblog-neural-probabilistic-logic-programming
AUTHORS: Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, Luc De Raedt
HIGHLIGHT: We introduce DeepProbLog, a probabilistic logic programming language that incorporates deep learning by means of neural predicates.

347, TITLE: Convergence of Cubic Regularization for Nonconvex Optimization under KL Property
https://papers.nips.cc/paper/7633-convergence-of-cubic-regularization-for-nonconvex-optimization-under-kl-property
AUTHORS: Yi Zhou, Zhe Wang, Yingbin Liang
HIGHLIGHT: In this paper, we explore the asymptotic convergence rate of CR by exploiting the ubiquitous Kurdyka-Lojasiewicz (KL) property of the nonconvex objective functions.

348, TITLE: Direct Estimation of Differences in Causal Graphs
https://papers.nips.cc/paper/7634-direct-estimation-of-differences-in-causal-graphs
AUTHORS: Yuhao Wang, Chandler Squires, Anastasiya Belyaeva, Caroline Uhler
HIGHLIGHT: We consider the problem of estimating the differences between two causal directed acyclic graph (DAG) models with a shared topological order given i.i.d. samples from each model.

349, TITLE: Sublinear Time Low-Rank Approximation of Distance Matrices
https://papers.nips.cc/paper/7635-sublinear-time-low-rank-approximation-of-distance-matrices
AUTHORS: Ainesh Bakshi, David Woodruff
HIGHLIGHT: We develop a recursive algorithm based on additive projection-cost preserving sampling.

350, TITLE: Variational PDEs for Acceleration on Manifolds and Application to Diffeomorphisms
https://papers.nips.cc/paper/7636-variational-pdes-for-acceleration-on-manifolds-and-application-to-diffeomorphisms
AUTHORS: Ganesh Sundaramoorthi, Anthony Yezzi
HIGHLIGHT: We consider the optimization of cost functionals on manifolds and derive a variational approach to accelerated methods on manifolds.

351, TITLE: Bayesian Inference of Temporal Task Specifications from Demonstrations
https://papers.nips.cc/paper/7637-bayesian-inference-of-temporal-task-specifications-from-demonstrations
AUTHORS: Ankit Shah, Pritish Kamath, Julie A. Shah, Shen Li
HIGHLIGHT: Inspired by this, we present Bayesian specification inference, a probabilistic model for inferring task specification as a temporal logic formula.

352, TITLE: Data center cooling using model-predictive control
https://papers.nips.cc/paper/7638-data-center-cooling-using-model-predictive-control
AUTHORS: Nevena Lazic, Craig Boutilier, Tyler Lu, Eehern Wong, Binz Roy, MK Ryu, Greg Imwalle
HIGHLIGHT: In this paper, we describe an application of RL “in the wild” to the task of regulating temperatures and airflow inside a large-scale data center (DC).

353, TITLE: Acceleration through Optimistic No-Regret Dynamics
https://papers.nips.cc/paper/7639-acceleration-through-optimistic-no-regret-dynamics
AUTHORS: Jun-Kun Wang, Jacob D. Abernethy
HIGHLIGHT: In this paper we show that the technique can be enhanced to a rate of $O(1/T^2)$ by extending recent work \cite{RS13,SALS15} that leverages \textit{optimistic learning} to speed up equilibrium computation.

354, TITLE: Lipschitz regularity of deep neural networks: analysis and efficient estimation
https://papers.nips.cc/paper/7640-lipschitz-regularity-of-deep-neural-networks-analysis-and-efficient-estimation
AUTHORS: Aladin Virmaux, Kevin Scaman
HIGHLIGHT: In this paper, we investigate one of the key characteristics to assess the regularity of such methods: the Lipschitz constant of deep learning architectures.

355, TITLE: Minimax Estimation of Neural Net Distance
https://papers.nips.cc/paper/7641-minimax-estimation-of-neural-net-distance
AUTHORS: Kaiyi Ji, Yingbin Liang
HIGHLIGHT: This paper investigates the minimax estimation problem of the neural net distance based on samples drawn from the distributions.

356, TITLE: Leveraging the Exact Likelihood of Deep Latent Variable Models
https://papers.nips.cc/paper/7642-leveraging-the-exact-likelihood-of-deep-latent-variable-models
AUTHORS: Pierre-Alexandre Mattei, Jes Frellsen
HIGHLIGHT: The purpose of this work is to study the general properties of this quantity and to show how they can be leveraged in practice.

357, TITLE: Bipartite Stochastic Block Models with Tiny Clusters
https://papers.nips.cc/paper/7643-bipartite-stochastic-block-models-with-tiny-clusters
AUTHORS: Stefan Neumann
HIGHLIGHT: We present a simple two-step algorithm which provably finds even tiny clusters of size $O(n^\epsilon)$, where n is the number of vertices in the graph and $\epsilon > 0$.

358, TITLE: Learning sparse neural networks via sensitivity-driven regularization
https://papers.nips.cc/paper/7644-learning-sparse-neural-networks-via-sensitivity-driven-regularization
AUTHORS: Enzo Tartaglione, Skjalg Leps?y, Attilio Fiandrotti, Gianluca Francini
HIGHLIGHT: In this context we quantify the output sensitivity to the parameters (i.e. their relevance to the network output) and introduce a regularization term that gradually lowers the absolute value of parameters with low sensitivity.

359, TITLE: Faster Online Learning of Optimal Threshold for Consistent F-measure Optimization
https://papers.nips.cc/paper/7645-faster-online-learning-of-optimal-threshold-for-consistent-f-measure-optimization
AUTHORS: Xiaoxuan Zhang, Mingrui Liu, Xun Zhou, Tianbao Yang
HIGHLIGHT: In this paper, we consider online F-measure optimization (OFO).

360, TITLE: Direct Runge-Kutta Discretization Achieves Acceleration
https://papers.nips.cc/paper/7646-direct-runge-kutta-discretization-achieves-acceleration
AUTHORS: Jingzhao Zhang, Aryan Mokhtari, Suvrit Sra, Ali Jadbabaie
HIGHLIGHT: We study gradient-based optimization methods obtained by directly discretizing a second-order ordinary differential equation (ODE) related to the continuous limit of Nesterov's accelerated gradient method.

361, TITLE: Adversarial Examples that Fool both Computer Vision and Time-Limited Humans
https://papers.nips.cc/paper/7647-adversarial-examples-that-fool-both-computer-vision-and-time-limited-humans
AUTHORS: Gamaleldin Elsayed, Shreya Shankar, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Goodfellow, Jascha Sohl-Dickstein
HIGHLIGHT: We find that adversarial examples that strongly transfer across computer vision models influence the classifications made by time-limited human observers.

362, TITLE: Stochastic Nested Variance Reduced Gradient Descent for Nonconvex Optimization
https://papers.nips.cc/paper/7648-stochastic-nested-variance-reduced-gradient-descent-for-nonconvex-optimization
AUTHORS: Dongruo Zhou, Pan Xu, Quanquan Gu
HIGHLIGHT: We propose a new stochastic gradient descent algorithm based on nested variance reduction.

363, TITLE: Faster Neural Networks Straight from JPEG
https://papers.nips.cc/paper/7649-faster-neural-networks-straight-from-jpeg
AUTHORS: Lionel Gueguen, Alex Sergeev, Ben Kadlec, Rosanne Liu, Jason Yosinski
HIGHLIGHT: In this paper we propose and explore a simple idea: train CNNs directly on the blockwise discrete cosine transform (DCT) coefficients computed and available in the middle of the JPEG codec.

364, TITLE: TopRank: A practical algorithm for online stochastic ranking
https://papers.nips.cc/paper/7650-toprank-a-practical-algorithm-for-online-stochastic-ranking
AUTHORS: Tor Lattimore, Branislav Kveton, Shuai Li, Csaba Szepesvari
HIGHLIGHT: We propose a generalized click model that encompasses many existing models, including the position-based and cascade models.

365, TITLE: Learning from discriminative feature feedback
https://papers.nips.cc/paper/7651-learning-from-discriminative-feature-feedback
AUTHORS: Sanjoy Dasgupta, Akansha Dey, Nicholas Roberts, Sivan Sabato
HIGHLIGHT: We present an efficient online algorithm for learning from such feedback and we give tight bounds on the number of mistakes made during the learning process.

366, TITLE: RetGK: Graph Kernels based on Return Probabilities of Random Walks

https://papers.nips.cc/paper/7652-retgk-graph-kernels-based-on-return-probabilities-of-random-walks
AUTHORS: Zhen Zhang, Mianzhi Wang, Yijian Xiang, Yan Huang, Arye Nehorai
HIGHLIGHT: In this paper, we develop a framework for computing graph kernels, based on return probabilities of random walks.

367, TITLE: Deep Generative Markov State Models
https://papers.nips.cc/paper/7653-deep-generative-markov-state-models
AUTHORS: Hao Wu, Andreas Mardt, Luca Pasquali, Frank Noe
HIGHLIGHT: We propose a deep generative Markov State Model (DeepGenMSM) learning framework for inference of metastable dynamical systems and prediction of trajectories.

368, TITLE: Early Stopping for Nonparametric Testing
https://papers.nips.cc/paper/7654-early-stopping-for-nonparametric-testing
AUTHORS: Meimei Liu, Guang Cheng
HIGHLIGHT: In this paper, we show that early stopping can also be applied to obtain the minimax optimal testing in a general non-parametric setup.

369, TITLE: Solving Non-smooth Constrained Programs with Lower Complexity than <var>\mathcal{O}(1/\varepsilon)</var>: A Primal-Dual Homotopy Smoothing Approach
https://papers.nips.cc/paper/7655-solving-non-smooth-constrained-programs-with-lower-complexity-than-mathcalo1varepsilon-a-primal-dual-homotopy-smoothing-approach
AUTHORS: Xiaohan Wei, Hao Yu, Qing Ling, Michael Neely
HIGHLIGHT: We propose a new primal-dual homotopy smoothing algorithm for a linearly constrained convex program, where neither the primal nor the dual function has to be smooth or strongly convex.

370, TITLE: Heterogeneous Bitwidth Binarization in Convolutional Neural Networks
https://papers.nips.cc/paper/7656-heterogeneous-bitwidth-binarization-in-convolutional-neural-networks
AUTHORS: Joshua Fromm, Shwetak Patel, Matthai Philipose
HIGHLIGHT: In this paper, we show that it is feasible and useful to select bitwidths at the parameter granularity during training.

371, TITLE: Unsupervised Learning of Object Landmarks through Conditional Image Generation
https://papers.nips.cc/paper/7657-unsupervised-learning-of-object-landmarks-through-conditional-image-generation
AUTHORS: Tomas Jakab, Ankush Gupta, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: We propose a method for learning landmark detectors for visual objects (such as the eyes and the nose in a face) without any manual supervision.

372, TITLE: Probabilistic Neural Programmed Networks for Scene Generation
https://papers.nips.cc/paper/7658-probabilistic-neural-programmed-networks-for-scene-generation
AUTHORS: Zhiwei Deng, Jiacheng Chen, YIFANG FU, Greg Mori
HIGHLIGHT: In this paper we address the text to scene image generation problem.

373, TITLE: The streaming rollout of deep networks - towards fully model-parallel execution
https://papers.nips.cc/paper/7659-the-streaming-rollout-of-deep-networks-towards-fully-model-parallel-execution
AUTHORS: Volker Fischer, Jan Koehler, Thomas Pfeil
HIGHLIGHT: In this study, we present a theoretical framework to describe rollouts, the level of model-parallelization they induce, and demonstrate differences in solving specific tasks.

374, TITLE: KONG: Kernels for ordered-neighborhood graphs
https://papers.nips.cc/paper/7660-kong-kernels-for-ordered-neighborhood-graphs
AUTHORS: Moez Draief, Konstantin Kutzkov, Kevin Scaman, Milan Vojnovic
HIGHLIGHT: We present novel graph kernels for graphs with node and edge labels that have ordered neighborhoods, i.e. when neighbor nodes follow an order.

375, TITLE: GumBolt: Extending Gumbel trick to Boltzmann priors
https://papers.nips.cc/paper/7661-gumbolt-extending-gumbel-trick-to-boltzmann-priors
AUTHORS: Amir H. Khoshaman, Mohammad Amin
HIGHLIGHT: Here, we propose the GumBolt, a model that extends the Gumbel trick to BM priors in VAEs.

376, TITLE: Neural Networks Trained to Solve Differential Equations Learn General Representations

https://papers.nips.cc/paper/7662-neural-networks-trained-to-solve-differential-equations-learn-general-representations
AUTHORS: Martin Magill, Faisal Qureshi, Hendrick de Haan
HIGHLIGHT: We introduce a technique based on the singular vector canonical correlation analysis (SVCCA) for measuring the generality of neural network layers across a continuously-parametrized set of tasks.

377, TITLE: Beauty-in-averageness and its contextual modulations: A Bayesian statistical account
https://papers.nips.cc/paper/7663-beauty-in-averageness-and-its-contextual-modulations-a-bayesian-statistical-account
AUTHORS: Chaitanya Ryali, Angela J. Yu
HIGHLIGHT: Existing models generally assume these preferences to be fixed.

378, TITLE: Distributed Weight Consolidation: A Brain Segmentation Case Study
https://papers.nips.cc/paper/7664-distributed-weight-consolidation-a-brain-segmentation-case-study
AUTHORS: Patrick McClure, Charles Y. Zheng, Jakub Kaczmarzyk, John Rogers-Lee, Satra Ghosh, Dylan Nielson, Peter A. Bandettini, Francisco Pereira
HIGHLIGHT: In this paper, we introduce distributed weight consolidation (DWC), a continual learning method to consolidate the weights of separate neural networks, each trained on an independent dataset.

379, TITLE: Efficient Projection onto the Perfect Phylogeny Model
https://papers.nips.cc/paper/7665-efficient-projection-onto-the-perfect-phylogeny-model
AUTHORS: Bei Jia, Surjyendu Ray, Sam Safavi, Jos? Bento
HIGHLIGHT: In this paper, we use Moreau's decomposition for proximal operators, and a tree reduction scheme, to develop a new algorithm to compute this projection.

380, TITLE: TETRIS: TilE-matching the TRemendous Irregular Sparsity
https://papers.nips.cc/paper/7666-tetris-tile-matching-the-tremendous-irregular-sparsity
AUTHORS: Yu Ji, Ling Liang, Lei Deng, Youyang Zhang, Youhui Zhang, Yuan Xie
HIGHLIGHT: In this work, we propose a novel method, TETRIS, to achieve both better hardware utilization and higher sparsity.

381, TITLE: Cooperative neural networks (CoNN): Exploiting prior independence structure for improved classification
https://papers.nips.cc/paper/7667-cooperative-neural-networks-conn-exploiting-prior-independence-structure-for-improved-classification
AUTHORS: Harsh Shrivastava, Eugene Bart, Bob Price, Hanjun Dai, Bo Dai, Srinivas Aluru
HIGHLIGHT: We propose a new approach, called cooperative neural networks (CoNN), which use a set of cooperatively trained neural networks to capture latent representations that exploit prior given independence structure.

382, TITLE: Differentially Private Robust Low-Rank Approximation
https://papers.nips.cc/paper/7668-differentially-private-robust-low-rank-approximation
AUTHORS: Raman Arora, Vladimir braverman, Jalaj Upadhyay
HIGHLIGHT: In this paper, we study the following robust low-rank matrix approximation problem: given a matrix $A \in \R^{n \times d}$, find a rank-k matrix B, while satisfying differential privacy, such that $ \norm{ A - B }_p \leq \alpha \mathsf{OPT}_k(A) + \tau,$ where $\norm{ M }_p$ is the entry-wise ℓ_p-norm and $\mathsf{OPT}_k(A):=\min_{\mathsf{rank}(X) \leq k} \norm{ A - X}_p$.

383, TITLE: Meta-Learning MCMC Proposals
https://papers.nips.cc/paper/7669-meta-learning-mcmc-proposals
AUTHORS: Tongzhou Wang, YI WU, Dave Moore, Stuart J. Russell
HIGHLIGHT: Inspired by recent progresses in meta-learning for training learning agents that can generalize to unseen environments, we propose a meta-learning approach to building effective and generalizable MCMC proposals.

384, TITLE: An Information-Theoretic Analysis for Thompson Sampling with Many Actions
https://papers.nips.cc/paper/7670-an-information-theoretic-analysis-for-thompson-sampling-with-many-actions
AUTHORS: Shi Dong, Benjamin Van Roy
HIGHLIGHT: We establish new bounds that depend instead on a notion of rate-distortion.

385, TITLE: Flexible and accurate inference and learning for deep generative models
https://papers.nips.cc/paper/7671-flexible-and-accurate-inference-and-learning-for-deep-generative-models
AUTHORS: Eszter V?rtes, Maneesh Sahani
HIGHLIGHT: We introduce a new approach to learning in hierarchical latent-variable generative models called the “distributed distributional code Helmholtz machine”, which emphasises flexibility and accuracy in the inferential process.

386, TITLE: The Price of Privacy for Low-rank Factorization
https://papers.nips.cc/paper/7672-the-price-of-privacy-for-low-rank-factorization
AUTHORS: Jalaj Upadhyay
HIGHLIGHT: In this paper, we study what price one has to pay to release \emph{differentially private low-rank factorization} of a matrix.
We present the first set of differentially private algorithms for all these settings.

387, TITLE: Regret Bounds for Robust Adaptive Control of the Linear Quadratic Regulator
https://papers.nips.cc/paper/7673-regret-bounds-for-robust-adaptive-control-of-the-linear-quadratic-regulator
AUTHORS: Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, Stephen Tu
HIGHLIGHT: Leveraging recent developments in the estimation of linear systems and in robust controller synthesis, we present the first provably polynomial time algorithm that achieves sub-linear regret on this problem.

388, TITLE: Bilevel Distance Metric Learning for Robust Image Recognition
https://papers.nips.cc/paper/7674-bilevel-distance-metric-learning-for-robust-image-recognition
AUTHORS: Jie Xu, Lei Luo, Cheng Deng, Heng Huang
HIGHLIGHT: In this paper, we integrate both feature extraction and metric learning into one joint optimization framework and propose a new bilevel distance metric learning model.

389, TITLE: Differentially Private Uniformly Most Powerful Tests for Binomial Data
https://papers.nips.cc/paper/7675-differentially-private-uniformly-most-powerful-tests-for-binomial-data
AUTHORS: Jordan Awan, Aleksandra Slavkovic
HIGHLIGHT: We derive uniformly most powerful (UMP) tests for simple and one-sided hypotheses for a population proportion within the framework of Differential Privacy (DP), optimizing finite sample performance.

390, TITLE: Scalable Coordinated Exploration in Concurrent Reinforcement Learning
https://papers.nips.cc/paper/7676-scalable-coordinated-exploration-in-concurrent-reinforcement-learning
AUTHORS: Maria Dimakopoulou, Ian Osband, Benjamin Van Roy
HIGHLIGHT: We consider a team of reinforcement learning agents that concurrently operate in a common environment, and we develop an approach to efficient coordinated exploration that is suitable for problems of practical scale.

391, TITLE: Integrated accounts of behavioral and neuroimaging data using flexible recurrent neural network models
https://papers.nips.cc/paper/7677-integrated-accounts-of-behavioral-and-neuroimaging-data-using-flexible-recurrent-neural-network-models
AUTHORS: Amir Dezfouli, Richard Morris, Fabio T. Ramos, Peter Dayan, Bernard Balleine
HIGHLIGHT: To address this limitation, we introduce a new method using recurrent neural network models that are flexible enough to be jointly fitted to the behavioral and neural data.

392, TITLE: BML: A High-performance, Low-cost Gradient Synchronization Algorithm for DML Training
https://papers.nips.cc/paper/7678-bml-a-high-performance-low-cost-gradient-synchronization-algorithm-for-dml-training
AUTHORS: Songtao Wang, Dan Li, Yang Cheng, Jinkun Geng, Yanshu Wang, Shuai Wang, Shu-Tao Xia, Jianping Wu
HIGHLIGHT: In this paper we propose BML, a new gradient synchronization algorithm with higher network performance and lower network cost than the current practice.

393, TITLE: Inexact trust-region algorithms on Riemannian manifolds
https://papers.nips.cc/paper/7679-inexact-trust-region-algorithms-on-riemannian-manifolds
AUTHORS: Hiroyuki Kasai, Bamdev Mishra
HIGHLIGHT: Addressing large-scale finite-sum problems, we specifically propose sub-sampled algorithms with a fixed bound on sub-sampled Hessian and gradient sizes, where the gradient and Hessian are computed by a random sampling technique.

394, TITLE: Can We Gain More from Orthogonality Regularizations in Training Deep Networks?
https://papers.nips.cc/paper/7680-can-we-gain-more-from-orthogonality-regularizations-in-training-deep-networks
AUTHORS: Nitin Bansal, Xiaohan Chen, Zhangyang Wang
HIGHLIGHT: This paper seeks to answer the question: as the (near-) orthogonality of weights is found to be a favorable property for training deep convolutional neural networks, how can we enforce it in more effective and easy-to-use ways?

395, TITLE: Binary Rating Estimation with Graph Side Information
https://papers.nips.cc/paper/7681-binary-rating-estimation-with-graph-side-information
AUTHORS: Kwangjun Ahn, Kangwook Lee, Hyunseung Cha, Changho Suh

HIGHLIGHT: In this work, we study the binary rating estimation problem to understand the fundamental value of graph side information.

396, TITLE: SimplE Embedding for Link Prediction in Knowledge Graphs
https://papers.nips.cc/paper/7682-simple-embedding-for-link-prediction-in-knowledge-graphs
AUTHORS: Seyed Mehran Kazemi, David Poole
HIGHLIGHT: Link prediction approaches aim at predicting new links for a knowledge graph given the existing links among the entities.

397, TITLE: Differentially Private Contextual Linear Bandits
https://papers.nips.cc/paper/7683-differentially-private-contextual-linear-bandits
AUTHORS: Roshan Shariff, Or Sheffet
HIGHLIGHT: Our goal is to devise private learners for the contextual linear bandit problem.

398, TITLE: Submodular Field Grammars: Representation, Inference, and Application to Image Parsing
https://papers.nips.cc/paper/7684-submodular-field-grammars-representation-inference-and-application-to-image-parsing
AUTHORS: Abram L. Friesen, Pedro M. Domingos
HIGHLIGHT: In this paper, we address this problem by associating with each production a submodular Markov random field whose labels are the subparts and whose labeling segments the current object into these subparts.

399, TITLE: A Bridging Framework for Model Optimization and Deep Propagation
https://papers.nips.cc/paper/7685-a-bridging-framework-for-model-optimization-and-deep-propagation
AUTHORS: Risheng Liu, Shichao Cheng, xiaokun liu, Long Ma, Xin Fan, Zhongxuan Luo
HIGHLIGHT: In this work, we provide a new paradigm, named Propagation and Optimization based Deep Model (PODM), to bridge the gaps between these different mechanisms (i.e., model optimization and deep propagation).

400, TITLE: Completing State Representations using Spectral Learning
https://papers.nips.cc/paper/7686-completing-state-representations-using-spectral-learning
AUTHORS: Nan Jiang, Alex Kulesza, Satinder Singh
HIGHLIGHT: In this paper we develop a novel algorithm for incorporating domain knowledge, in the form of an imperfect state representation, as side information to speed spectral learning for PSRs.

401, TITLE: Optimization of Smooth Functions with Noisy Observations: Local Minimax Rates
https://papers.nips.cc/paper/7687-optimization-of-smooth-functions-with-noisy-observations-local-minimax-rates
AUTHORS: Yining Wang, Sivaraman Balakrishnan, Aarti Singh
HIGHLIGHT: We propose a local minimax framework to study the fundamental difficulty of optimizing smooth functions with adaptive function evaluations.

402, TITLE: Adding One Neuron Can Eliminate All Bad Local Minima
https://papers.nips.cc/paper/7688-adding-one-neuron-can-eliminate-all-bad-local-minima
AUTHORS: SHIYU LIANG, Ruoyu Sun, Jason D. Lee, R. Srikant
HIGHLIGHT: In this paper, we study the landscape of neural networks for binary classification tasks.

403, TITLE: Mean-field theory of graph neural networks in graph partitioning
https://papers.nips.cc/paper/7689-mean-field-theory-of-graph-neural-networks-in-graph-partitioning
AUTHORS: Tatsuro Kawamoto, Masashi Tsubaki, Tomoyuki Obuchi
HIGHLIGHT: A theoretical performance analysis of the graph neural network (GNN) is presented.

404, TITLE: The Physical Systems Behind Optimization Algorithms
https://papers.nips.cc/paper/7690-the-physical-systems-behind-optimization-algorithms
AUTHORS: Lin Yang, Raman Arora, Vladimir braverman, Tuo Zhao
HIGHLIGHT: In particular, we study gradient descent, proximal gradient descent, coordinate gradient descent, proximal coordinate gradient, and Newton's methods as well as their Nesterov's accelerated variants in a unified framework motivated by a natural connection of optimization algorithms to physical systems.

405, TITLE: Mallows Models for Top-k Lists
https://papers.nips.cc/paper/7691-mallows-models-for-top-k-lists
AUTHORS: Flavio Chierichetti, Anirban Dasgupta, Shahrzad Haddadan, Ravi Kumar, Silvio Lattanzi
HIGHLIGHT: Motivated by common practical situations, in this paper, we generalize Mallows to model distributions on top-k lists by using a suitable distance measure between top-k lists.

406, TITLE: Amortized Inference Regularization
https://papers.nips.cc/paper/7692-amortized-inference-regularization
AUTHORS: Rui Shu, Hung H. Bui, Shengjia Zhao, Mykel J. Kochenderfer, Stefano Ermon
HIGHLIGHT: In this paper, we leverage the fact that VAEs rely on amortized inference and propose techniques for amortized inference regularization (AIR) that control the smoothness of the inference model.

407, TITLE: Maximum Causal Tsallis Entropy Imitation Learning
https://papers.nips.cc/paper/7693-maximum-causal-tsallis-entropy-imitation-learning
AUTHORS: Kyungjae Lee, Sungjoon Choi, Songhwai Oh
HIGHLIGHT: In this paper, we propose a novel maximum causal Tsallis entropy (MCTE) framework for imitation learning which can efficiently learn a sparse multi-modal policy distribution from demonstrations.

408, TITLE: Limited Memory Kelley's Method Converges for Composite Convex and Submodular Objectives
https://papers.nips.cc/paper/7694-limited-memory-kelleys-method-converges-for-composite-convex-and-submodular-objectives
AUTHORS: Song Zhou, Swati Gupta, Madeleine Udell
HIGHLIGHT: We propose a limited memory version of Kelley's method (L-KM) and of OSM that requires limited memory (at most n+ 1 constraints for an n-dimensional problem) independent of the iteration.

409, TITLE: Semi-Supervised Learning with Declaratively Specified Entropy Constraints
https://papers.nips.cc/paper/7695-semi-supervised-learning-with-declaratively-specified-entropy-constraints
AUTHORS: Haitian Sun, William W. Cohen, Lidong Bing
HIGHLIGHT: We propose a technique for declaratively specifying strategies for semi-supervised learning (SSL).

410, TITLE: End-to-end Symmetry Preserving Inter-atomic Potential Energy Model for Finite and Extended Systems
https://papers.nips.cc/paper/7696-end-to-end-symmetry-preserving-inter-atomic-potential-energy-model-for-finite-and-extended-systems
AUTHORS: Linfeng Zhang, Jiequn Han, Han Wang, Wissam Saidi, Roberto Car, Weinan E
HIGHLIGHT: Here we develop Deep Potential - Smooth Edition (DeepPot-SE), an end-to-end machine learning-based PES model, which is able to efficiently represent the PES for a wide variety of systems with the accuracy of ab initio quantum mechanics models.

411, TITLE: Sparsified SGD with Memory
https://papers.nips.cc/paper/7697-sparsified-sgd-with-memory
AUTHORS: Sebastian U. Stich, Jean-Baptiste Cordonnier, Martin Jaggi
HIGHLIGHT: In this work we analyze Stochastic Gradient Descent (SGD) with k-sparsification or compression (for instance top-k or random-k) and show that this scheme converges at the same rate as vanilla SGD when equipped with error compensation (keeping track of accumulated errors in memory).

412, TITLE: Exponentiated Strongly Rayleigh Distributions
https://papers.nips.cc/paper/7698-exponentiated-strongly-rayleigh-distributions
AUTHORS: Zelda E. Mariet, Suvrit Sra, Stefanie Jegelka
HIGHLIGHT: We introduce in this paper Exponentiated Strongly Rayleigh (ESR) measures, which sharpen (or smoothen) the negative dependence property of SR measures via a single parameter (the exponent) that can intuitively understood as an inverse temperature.

413, TITLE: Importance Weighting and Variational Inference
https://papers.nips.cc/paper/7699-importance-weighting-and-variational-inference
AUTHORS: Justin Domke, Daniel R. Sheldon
HIGHLIGHT: Recent work used importance sampling ideas for better variational bounds on likelihoods.

414, TITLE: Transfer Learning from Speaker Verification to Multispeaker Text-To-Speech Synthesis
https://papers.nips.cc/paper/7700-transfer-learning-from-speaker-verification-to-multispeaker-text-to-speech-synthesis
AUTHORS: Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, zhifeng Chen, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu
HIGHLIGHT: We describe a neural network-based system for text-to-speech (TTS) synthesis that is able to generate speech audio in the voice of many different speakers, including those unseen during training.

415, TITLE: Expanding Holographic Embeddings for Knowledge Completion
https://papers.nips.cc/paper/7701-expanding-holographic-embeddings-for-knowledge-completion

AUTHORS: Yexiang Xue, Yang Yuan, Zhitian Xu, Ashish Sabharwal
HIGHLIGHT: We propose a new family of embeddings for knowledge graphs that interpolate between a method with high model complexity and one, namely Holographic embeddings (HolE), with low dimensionality and high training efficiency.

416, TITLE: Lifelong Inverse Reinforcement Learning
https://papers.nips.cc/paper/7702-lifelong-inverse-reinforcement-learning
AUTHORS: Jorge Armando Mendez Mendez, Shashank Shivkumar, Eric Eaton
HIGHLIGHT: As one solution to this problem, we propose the first lifelong learning approach to inverse reinforcement learning, which learns consecutive tasks via demonstration, continually transferring knowledge between tasks to improve performance.
To address this challenge, we introduce the novel problem of lifelong learning from demonstration, which allows the agent to continually build upon knowledge learned from previously demonstrated tasks to accelerate the learning of new tasks, reducing the amount of demonstrations required.

417, TITLE: Explaining Deep Learning Models -- A Bayesian Non-parametric Approach
https://papers.nips.cc/paper/7703-explaining-deep-learning-models-a-bayesian-non-parametric-approach
AUTHORS: Wenbo Guo, Sui Huang, Yunzhe Tao, Xinyu Xing, Lin Lin
HIGHLIGHT: In this work, we propose a novel technical approach that augments a Bayesian non-parametric regression mixture model with multiple elastic nets.

418, TITLE: Third-order Smoothness Helps: Faster Stochastic Optimization Algorithms for Finding Local Minima
https://papers.nips.cc/paper/7704-third-order-smoothness-helps-faster-stochastic-optimization-algorithms-for-finding-local-minima
AUTHORS: Yaodong Yu, Pan Xu, Quanquan Gu
HIGHLIGHT: We propose stochastic optimization algorithms that can find local minima faster than existing algorithms for nonconvex optimization problems, by exploiting the third-order smoothness to escape non-degenerate saddle points more efficiently.

419, TITLE: COLA: Decentralized Linear Learning
https://papers.nips.cc/paper/7705-cola-decentralized-linear-learning
AUTHORS: Lie He, An Bian, Martin Jaggi
HIGHLIGHT: We propose COLA, a new decentralized training algorithm with strong theoretical guarantees and superior practical performance.

420, TITLE: MiME: Multilevel Medical Embedding of Electronic Health Records for Predictive Healthcare
https://papers.nips.cc/paper/7706-mime-multilevel-medical-embedding-of-electronic-health-records-for-predictive-healthcare
AUTHORS: Edward Choi, Cao Xiao, Walter Stewart, Jimeng Sun
HIGHLIGHT: We propose Multilevel Medical Embedding (MiME) which learns the multilevel embedding of EHR data while jointly performing auxiliary prediction tasks that rely on this inherent EHR structure without the need for external labels.

421, TITLE: Adaptive Sampling Towards Fast Graph Representation Learning
https://papers.nips.cc/paper/7707-adaptive-sampling-towards-fast-graph-representation-learning
AUTHORS: Wenbing Huang, Tong Zhang, Yu Rong, Junzhou Huang
HIGHLIGHT: In this paper, we accelerate the training of GCNs through developing an adaptive layer-wise sampling method.

422, TITLE: Hunting for Discriminatory Proxies in Linear Regression Models
https://papers.nips.cc/paper/7708-hunting-for-discriminatory-proxies-in-linear-regression-models
AUTHORS: Samuel Yeom, Anupam Datta, Matt Fredrikson
HIGHLIGHT: In this paper we formulate a definition of proxy use for the setting of linear regression and present algorithms for detecting proxies.

423, TITLE: Towards Robust Detection of Adversarial Examples
https://papers.nips.cc/paper/7709-towards-robust-detection-of-adversarial-examples
AUTHORS: Tianyu Pang, Chao Du, Yinpeng Dong, Jun Zhu
HIGHLIGHT: In this paper, we present a novel training procedure and a thresholding test strategy, towards robust detection of adversarial examples.

424, TITLE: Active Matting
https://papers.nips.cc/paper/7710-active-matting
AUTHORS: Xin Yang, Ke Xu, Shaozhe Chen, Shengfeng He, Baocai Yin Yin, Rynson Lau
HIGHLIGHT: In this paper, we explore the intrinsic relationship between the user input and the matting algorithm to address the problem of where and when the user should provide the input.

425, TITLE: Learning filter widths of spectral decompositions with wavelets
https://papers.nips.cc/paper/7711-learning-filter-widths-of-spectral-decompositions-with-wavelets
AUTHORS: Haidar Khan, Bulent Yener
HIGHLIGHT: We propose the wavelet deconvolution (WD) layer as an efficient alternative to this preprocessing step that eliminates a significant number of hyperparameters.

426, TITLE: Byzantine Stochastic Gradient Descent
https://papers.nips.cc/paper/7712-byzantine-stochastic-gradient-descent
AUTHORS: Dan Alistarh, Zeyuan Allen-Zhu, Jerry Li
HIGHLIGHT: This paper studies the problem of distributed stochastic optimization in an adversarial setting where, out of m machines which allegedly compute stochastic gradients every iteration, an α-fraction are Byzantine, and may behave adversarially.

427, TITLE: PG-TS: Improved Thompson Sampling for Logistic Contextual Bandits
https://papers.nips.cc/paper/7713-pg-ts-improved-thompson-sampling-for-logistic-contextual-bandits
AUTHORS: Bianca Dumitrascu, Karen Feng, Barbara Engelhardt
HIGHLIGHT: Using a fast inference procedure with Polya-Gamma distributed augmentation variables, we propose an improved version of Thompson Sampling, a Bayesian formulation of contextual bandits with near-optimal performance.

428, TITLE: Spectral Filtering for General Linear Dynamical Systems
https://papers.nips.cc/paper/7714-spectral-filtering-for-general-linear-dynamical-systems
AUTHORS: Elad Hazan, HOLDEN LEE, Karan Singh, Cyril Zhang, Yi Zhang
HIGHLIGHT: We give a polynomial-time algorithm for learning latent-state linear dynamical systems without system identification, and without assumptions on the spectral radius of the system's transition matrix.

429, TITLE: On Learning Intrinsic Rewards for Policy Gradient Methods
https://papers.nips.cc/paper/7715-on-learning-intrinsic-rewards-for-policy-gradient-methods
AUTHORS: Zeyu Zheng, Junhyuk Oh, Satinder Singh
HIGHLIGHT: In this paper we build on the Optimal Rewards Framework of Singh et al. that defines the optimal intrinsic reward function as one that when used by an RL agent achieves behavior that optimizes the task-specifying or extrinsic reward function.

430, TITLE: Boolean Decision Rules via Column Generation
https://papers.nips.cc/paper/7716-boolean-decision-rules-via-column-generation
AUTHORS: Sanjeeb Dash, Oktay Gunluk, Dennis Wei
HIGHLIGHT: To handle large datasets, we propose an approximate CG algorithm using randomization.

431, TITLE: Adversarial Text Generation via Feature-Mover's Distance
https://papers.nips.cc/paper/7717-adversarial-text-generation-via-feature-movers-distance
AUTHORS: Liqun Chen, Shuyang Dai, Chenyang Tao, Haichao Zhang, Zhe Gan, Dinghan Shen, Yizhe Zhang, Guoyin Wang, Ruiyi Zhang, Lawrence Carin
HIGHLIGHT: Instead of using the standard GAN objective, we propose to improve text-generation GAN via a novel approach inspired by optimal transport.

432, TITLE: Fast Rates of ERM and Stochastic Approximation: Adaptive to Error Bound Conditions
https://papers.nips.cc/paper/7718-fast-rates-of-erm-and-stochastic-approximation-adaptive-to-error-bound-conditions
AUTHORS: Mingrui Liu, Xiaoxuan Zhang, Lijun Zhang, Jing Rong, Tianbao Yang
HIGHLIGHT: The main contributions of this paper are two-fold.

433, TITLE: Learning Bounds for Greedy Approximation with Explicit Feature Maps from Multiple Kernels
https://papers.nips.cc/paper/7719-learning-bounds-for-greedy-approximation-with-explicit-feature-maps-from-multiple-kernels
AUTHORS: Shahin Shahrampour, Vahid Tarokh
HIGHLIGHT: In this work, we tackle this problem by efficiently choosing such features from multiple kernels in a greedy fashion.

434, TITLE: A Mathematical Model For Optimal Decisions In A Representative Democracy
https://papers.nips.cc/paper/7720-a-mathematical-model-for-optimal-decisions-in-a-representative-democracy
AUTHORS: Malik Magdon-Ismail, Lirong Xia
HIGHLIGHT: We introduce a mathematical model for studying representative democracy, in particular understanding the parameters of a representative democracy that gives maximum decision making capability.

435, TITLE: Negotiable Reinforcement Learning for Pareto Optimal Sequential Decision-Making
https://papers.nips.cc/paper/7721-negotiable-reinforcement-learning-for-pareto-optimal-sequential-decision-making
AUTHORS: Nishant Desai, Andrew Critch, Stuart J. Russell
HIGHLIGHT: In this paper, we derive a more precise generalization for the sequential decision setting in the case of principals with different priors on the dynamics of the environment.

436, TITLE: Non-metric Similarity Graphs for Maximum Inner Product Search
https://papers.nips.cc/paper/7722-non-metric-similarity-graphs-for-maximum-inner-product-search
AUTHORS: Stanislav Morozov, Artem Babenko
HIGHLIGHT: In this paper we address the problem of Maximum Inner Product Search (MIPS) that is currently the computational bottleneck in a large number of machine learning applications.

437, TITLE: Recurrently Controlled Recurrent Networks
https://papers.nips.cc/paper/7723-recurrently-controlled-recurrent-networks
AUTHORS: Yi Tay, Anh Tuan Luu, Siu Cheung Hui
HIGHLIGHT: This paper proposes a recurrently controlled recurrent network (RCRN) for expressive and powerful sequence encoding.

438, TITLE: Fast greedy algorithms for dictionary selection with generalized sparsity constraints
https://papers.nips.cc/paper/7724-fast-greedy-algorithms-for-dictionary-selection-with-generalized-sparsity-constraints
AUTHORS: Kaito Fujii, Tasuku Soma
HIGHLIGHT: We propose a novel efficient greedy algorithm for dictionary selection.

439, TITLE: Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models
https://papers.nips.cc/paper/7725-deep-reinforcement-learning-in-a-handful-of-trials-using-probabilistic-dynamics-models
AUTHORS: Kurtland Chua, Roberto Calandra, Rowan McAllister, Sergey Levine
HIGHLIGHT: In this paper, we study how to bridge this gap, by employing uncertainty-aware dynamics models.

440, TITLE: A Smoother Way to Train Structured Prediction Models
https://papers.nips.cc/paper/7726-a-smoother-way-to-train-structured-prediction-models
AUTHORS: Venkata Krishna Pillutla, Vincent Roulet, Sham M. Kakade, Zaid Harchaoui
HIGHLIGHT: We present a framework to train a structured prediction model by performing smoothing on the inference algorithm it builds upon.

441, TITLE: Context-dependent upper-confidence bounds for directed exploration
https://papers.nips.cc/paper/7727-context-dependent-upper-confidence-bounds-for-directed-exploration
AUTHORS: Raksha Kumaraswamy, Matthew Schlegel, Adam White, Martha White
HIGHLIGHT: In this work, we provide a novel, computationally efficient, incremental exploration strategy, leveraging this property of least-squares temporal difference learning (LSTD).

442, TITLE: A Unified View of Piecewise Linear Neural Network Verification
https://papers.nips.cc/paper/7728-a-unified-view-of-piecewise-linear-neural-network-verification
AUTHORS: Rudy R. Bunel, Ilker Turkaslan, Philip Torr, Pushmeet Kohli, Pawan K. Mudigonda
HIGHLIGHT: To facilitate progress on this crucial area, we make two key contributions.
Second, we propose a new data set of benchmarks which includes a collection of previously released testcases.

443, TITLE: Hierarchical Graph Representation Learning with Differentiable Pooling
https://papers.nips.cc/paper/7729-hierarchical-graph-representation-learning-with-differentiable-pooling
AUTHORS: Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, Jure Leskovec
HIGHLIGHT: Here we propose DiffPool, a differentiable graph pooling module that can generate hierarchical representations of graphs and can be combined with various graph neural network architectures in an end-to-end fashion.

444, TITLE: Non-Ergodic Alternating Proximal Augmented Lagrangian Algorithms with Optimal Rates
https://papers.nips.cc/paper/7730-non-ergodic-alternating-proximal-augmented-lagrangian-algorithms-with-optimal-rates
AUTHORS: Quoc Tran Dinh
HIGHLIGHT: We develop two new non-ergodic alternating proximal augmented Lagrangian algorithms (NEAPAL) to solve a class of nonsmooth constrained convex optimization problems.

445, TITLE: Information-based Adaptive Stimulus Selection to Optimize Communication Efficiency in Brain-Computer Interfaces
https://papers.nips.cc/paper/7731-information-based-adaptive-stimulus-selection-to-optimize-communication-efficiency-in-brain-computer-interfaces
AUTHORS: Boyla Mainsah, Dmitry Kalika, Leslie Collins, Siyuan Liu, Chandra Throckmorton
HIGHLIGHT: We derive a simple analytical solution of an information-based objective function for BCI stimulus selection by transforming the high-dimensional stimulus space into a one-dimensional space that parameterizes the objective function - the prior probability mass of the stimulus under consideration, irrespective of its contents.

446, TITLE: Porcupine Neural Networks: Approximating Neural Network Landscapes
https://papers.nips.cc/paper/7732-porcupine-neural-networks-approximating-neural-network-landscapes
AUTHORS: Soheil Feizi, Hamid Javadi, Jesse Zhang, David Tse
HIGHLIGHT: In this paper, we take another approach to this problem by constraining the network such that the corresponding optimization landscape has good theoretical properties without significantly compromising performance.

447, TITLE: Fairness Through Computationally-Bounded Awareness
https://papers.nips.cc/paper/7733-fairness-through-computationally-bounded-awareness
AUTHORS: Michael Kim, Omer Reingold, Guy Rothblum
HIGHLIGHT: We study the problem of fair classification within the versatile framework of Dwork et al. [ITCS '12], which assumes the existence of a metric that measures similarity between pairs of individuals.

448, TITLE: Adaptive Negative Curvature Descent with Applications in Non-convex Optimization
https://papers.nips.cc/paper/7734-adaptive-negative-curvature-descent-with-applications-in-non-convex-optimization
AUTHORS: Mingrui Liu, Zhe Li, Xiaoyu Wang, Jinfeng Yi, Tianbao Yang
HIGHLIGHT: To address this issue, we propose an adaptive NCD to allow for an adaptive error dependent on the current gradient's magnitude in approximating the smallest eigen-value of the Hessian, and to encourage competition between a noisy NCD step and gradient descent step.

449, TITLE: Is Q-Learning Provably Efficient?
https://papers.nips.cc/paper/7735-is-q-learning-provably-efficient
AUTHORS: Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, Michael I. Jordan
HIGHLIGHT: We prove that, in an episodic MDP setting, Q-learning with UCB exploration achieves regret $\tlO(\sqrt{H^3 SAT})$ where S and A are the numbers of states and actions, H is the number of steps per episode, and T is the total number of steps.

450, TITLE: Interpreting Neural Network Judgments via Minimal, Stable, and Symbolic Corrections
https://papers.nips.cc/paper/7736-interpreting-neural-network-judgments-via-minimal-stable-and-symbolic-corrections
AUTHORS: Xin Zhang, Armando Solar-Lezama, Rishabh Singh
HIGHLIGHT: We present a new algorithm to generate minimal, stable, and symbolic corrections to an input that will cause a neural network with ReLU activations to change its output.

451, TITLE: Measures of distortion for machine learning
https://papers.nips.cc/paper/7737-measures-of-distortion-for-machine-learning
AUTHORS: Leena Chennuru Vankadara, Ulrike von Luxburg
HIGHLIGHT: In this paper, we show that many of the existing distortion measures behave in an undesired way, when considered from a machine learning point of view.

452, TITLE: On the Local Minima of the Empirical Risk
https://papers.nips.cc/paper/7738-on-the-local-minima-of-the-empirical-risk
AUTHORS: Chi Jin, Lydia T. Liu, Rong Ge, Michael I. Jordan
HIGHLIGHT: We propose a simple algorithm based on stochastic gradient descent (SGD) on a smoothed version of f that is guaranteed to achieve our goal as long as $\nu \le O(\epsilon^{1.5}/d)$.

453, TITLE: Densely Connected Attention Propagation for Reading Comprehension
https://papers.nips.cc/paper/7739-densely-connected-attention-propagation-for-reading-comprehension
AUTHORS: Yi Tay, Anh Tuan Luu, Siu Cheung Hui, Jian Su
HIGHLIGHT: We propose DecaProp (Densely Connected Attention Propagation), a new densely connected neural architecture for reading comprehension (RC).

454, TITLE: Bandit Learning with Positive Externalities
https://papers.nips.cc/paper/7740-bandit-learning-with-positive-externalities

AUTHORS: Virag Shah, Jose Blanchet, Ramesh Johari
HIGHLIGHT: We study multiarmed bandit (MAB) problems with positive externalities.

455, TITLE: Learning Confidence Sets using Support Vector Machines
https://papers.nips.cc/paper/7741-learning-confidence-sets-using-support-vector-machines
AUTHORS: Wenbo Wang, Xingye Qiao
HIGHLIGHT: Instead of plug-in approaches, we propose a support vector classifier to construct confidence sets in a flexible manner.

456, TITLE: Efficient Neural Network Robustness Certification with General Activation Functions
https://papers.nips.cc/paper/7742-efficient-neural-network-robustness-certification-with-general-activation-functions
AUTHORS: Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, Luca Daniel
HIGHLIGHT: To address this issue, in this paper we introduce CROWN, a general framework to certify robustness of neural networks with general activation functions.

457, TITLE: Hessian-based Analysis of Large Batch Training and Robustness to Adversaries
https://papers.nips.cc/paper/7743-hessian-based-analysis-of-large-batch-training-and-robustness-to-adversaries
AUTHORS: Zhewei Yao, Amir Gholami, Qi Lei, Kurt Keutzer, Michael W. Mahoney
HIGHLIGHT: We present detailed experiments with five different network architectures, including a residual network, tested on MNIST, CIFAR-10/100 datasets.

458, TITLE: Neural Edit Operations for Biological Sequences
https://papers.nips.cc/paper/7744-neural-edit-operations-for-biological-sequences
AUTHORS: Satoshi Koide, Keisuke Kawano, Takuro Kutsuna
HIGHLIGHT: Our analysis shows that CNNs can recognize star-free regular expressions, and that deeper CNNs can recognize more complex regular expressions including the insertion/deletion of characters.

459, TITLE: Objective and efficient inference for couplings in neuronal networks
https://papers.nips.cc/paper/7745-objective-and-efficient-inference-for-couplings-in-neuronal-networks
AUTHORS: Yu Terada, Tomoyuki Obuchi, Takuya Isomura, Yoshiyuki Kabashima
HIGHLIGHT: This significantly reduces the computational cost of the screening method employed in the proposed objective procedure, making it possible to treat large-size systems as in this study.

460, TITLE: Learning from Group Comparisons: Exploiting Higher Order Interactions
https://papers.nips.cc/paper/7746-learning-from-group-comparisons-exploiting-higher-order-interactions
AUTHORS: Yao Li, Minhao Cheng, Kevin Fujii, Fushing Hsieh, Cho-Jui Hsieh
HIGHLIGHT: In this paper, we propose a new model that takes the player-interaction effects into consideration.

461, TITLE: Supervising Unsupervised Learning
https://papers.nips.cc/paper/7747-supervising-unsupervised-learning
AUTHORS: Vikas Garg
HIGHLIGHT: We introduce a framework to transfer knowledge acquired from a repository of (heterogeneous) supervised datasets to new unsupervised datasets.

462, TITLE: Nonparametric Bayesian Lomax delegate racing for survival analysis with competing risks
https://papers.nips.cc/paper/7748-nonparametric-bayesian-lomax-delegate-racing-for-survival-analysis-with-competing-risks
AUTHORS: Quan Zhang, Mingyuan Zhou
HIGHLIGHT: We propose Lomax delegate racing (LDR) to explicitly model the mechanism of survival under competing risks and to interpret how the covariates accelerate or decelerate the time to event.

463, TITLE: Adversarially Robust Generalization Requires More Data
https://papers.nips.cc/paper/7749-adversarially-robust-generalization-requires-more-data
AUTHORS: Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, Aleksander Madry
HIGHLIGHT: To better understand this phenomenon, we study adversarially robust learning from the viewpoint of generalization.

464, TITLE: Improving Exploration in Evolution Strategies for Deep Reinforcement Learning via a Population of Novelty-Seeking Agents
https://papers.nips.cc/paper/7750-improving-exploration-in-evolution-strategies-for-deep-reinforcement-learning-via-a-population-of-novelty-seeking-agents

AUTHORS: Edoardo Conti, Vashisht Madhavan, Felipe Petroski Such, Joel Lehman, Kenneth Stanley, Jeff Clune
HIGHLIGHT: This paper thus introduces a family of fast, scalable algorithms for reinforcement learning that are capable of directed exploration.

465, TITLE: Practical exact algorithm for trembling-hand equilibrium refinements in games
https://papers.nips.cc/paper/7751-practical-exact-algorithm-for-trembling-hand-equilibrium-refinements-in-games
AUTHORS: Gabriele Farina, Nicola Gatti, Tuomas Sandholm
HIGHLIGHT: In this paper, we design an exact polynomial-time algorithm for finding trembling-hand equilibria in zero-sum extensive-form games.

466, TITLE: LAG: Lazily Aggregated Gradient for Communication-Efficient Distributed Learning
https://papers.nips.cc/paper/7752-lag-lazily-aggregated-gradient-for-communication-efficient-distributed-learning
AUTHORS: Tianyi Chen, Georgios Giannakis, Tao Sun, Wotao Yin
HIGHLIGHT: This paper presents a new class of gradient methods for distributed machine learning that adaptively skip the gradient calculations to learn with reduced communication and computation.

467, TITLE: Scalable Robust Matrix Factorization with Nonconvex Loss
https://papers.nips.cc/paper/7753-scalable-robust-matrix-factorization-with-nonconvex-loss
AUTHORS: Quanming Yao, James Kwok
HIGHLIGHT: In this paper, we propose the use of nonconvex loss to enhance robustness.

468, TITLE: Power-law efficient neural codes provide general link between perceptual bias and discriminability
https://papers.nips.cc/paper/7754-power-law-efficient-neural-codes-provide-general-link-between-perceptual-bias-and-discriminability
AUTHORS: Michael Morais, Jonathan W. Pillow
HIGHLIGHT: Here we generalize these results to show that the same law arises under a much larger family of optimal neural codes, introducing a unifying framework that we call power-law efficient coding.

469, TITLE: Geometry-Aware Recurrent Neural Networks for Active Visual Recognition
https://papers.nips.cc/paper/7755-geometry-aware-recurrent-neural-networks-for-active-visual-recognition
AUTHORS: Ricson Cheng, Ziyan Wang, Katerina Fragkiadaki
HIGHLIGHT: We present recurrent geometry-aware neural networks that integrate visual in- formation across multiple views of a scene into 3D latent feature tensors, while maintaining an one-to-one mapping between 3D physical locations in the world scene and latent feature locations.

470, TITLE: Unsupervised Adversarial Invariance
https://papers.nips.cc/paper/7756-unsupervised-adversarial-invariance
AUTHORS: Ayush Jaiswal, Rex Yue Wu, Wael Abd-Almageed, Prem Natarajan
HIGHLIGHT: We present a novel unsupervised invariance induction framework for neural networks that learns a split representation of data through competitive training between the prediction task and a reconstruction task coupled with disentanglement, without needing any labeled information about nuisance factors or domain knowledge.

471, TITLE: Content preserving text generation with attribute controls
https://papers.nips.cc/paper/7757-content-preserving-text-generation-with-attribute-controls
AUTHORS: Lajanugen Logeswaran, Honglak Lee, Samy Bengio
HIGHLIGHT: In this work, we address the problem of modifying textual attributes of sentences.

472, TITLE: Multi-armed Bandits with Compensation
https://papers.nips.cc/paper/7758-multi-armed-bandits-with-compensation
AUTHORS: Siwei Wang, Longbo Huang
HIGHLIGHT: We propose and study the known-compensation multi-arm bandit (KCMAB) problem, where a system controller offers a set of arms to many short-term players for T steps.

473, TITLE: GradiVeQ: Vector Quantization for Bandwidth-Efficient Gradient Aggregation in Distributed CNN Training
https://papers.nips.cc/paper/7759-gradiveq-vector-quantization-for-bandwidth-efficient-gradient-aggregation-in-distributed-cnn-training
AUTHORS: Mingchao Yu, Zhifeng Lin, Krishna Narra, Songze Li, Youjie Li, Nam Sung Kim, Alexander Schwing, Murali Annavaram, Salman Avestimehr
HIGHLIGHT: In this paper, we empirically demonstrate the strong linear correlations between CNN gradients, and propose a gradient vector quantization technique, named GradiVeQ, to exploit these correlations through principal component analysis (PCA) for substantial gradient dimension reduction.

474, TITLE: Learning in Games with Lossy Feedback
https://papers.nips.cc/paper/7760-learning-in-games-with-lossy-feedback
AUTHORS: Zhengyuan Zhou, Panayotis Mertikopoulos, Susan Athey, Nicholas Bambos, Peter W. Glynn, Yinyu Ye
HIGHLIGHT: We consider a game-theoretical multi-agent learning problem where the feedback information can be lost during the learning process and rewards are given by a broad class of games known as variationally stable games.

475, TITLE: Scalable methods for 8-bit training of neural networks
https://papers.nips.cc/paper/7761-scalable-methods-for-8-bit-training-of-neural-networks
AUTHORS: Ron Banner, Itay Hubara, Elad Hoffer, Daniel Soudry
HIGHLIGHT: Extensive research in the field suggests many different quantization schemes.

476, TITLE: Dropping Symmetry for Fast Symmetric Nonnegative Matrix Factorization
https://papers.nips.cc/paper/7762-dropping-symmetry-for-fast-symmetric-nonnegative-matrix-factorization
AUTHORS: Zhihui Zhu, Xiao Li, Kai Liu, Qiuwei Li
HIGHLIGHT: Unfortunately, designing fast algorithms for Symmetric NMF is not as easy as for the nonsymmetric counterpart, the latter admitting the splitting property that allows efficient alternating-type algorithms.

477, TITLE: Link Prediction Based on Graph Neural Networks
https://papers.nips.cc/paper/7763-link-prediction-based-on-graph-neural-networks
AUTHORS: Muhan Zhang, Yixin Chen
HIGHLIGHT: In this paper, we study this heuristic learning paradigm for link prediction.

478, TITLE: Why so gloomy? A Bayesian explanation of human pessimism bias in the multi-armed bandit task
https://papers.nips.cc/paper/7764-why-so-gloomy-a-bayesian-explanation-of-human-pessimism-bias-in-the-multi-armed-bandit-task
AUTHORS: Dalin Guo, Angela J. Yu
HIGHLIGHT: We present data from a human stationary bandit experiment, in which we vary the average abundance and variability of reward availability (mean and variance of reward rate distributions).

479, TITLE: Near-Optimal Time and Sample Complexities for Solving Markov Decision Processes with a Generative Model
https://papers.nips.cc/paper/7765-near-optimal-time-and-sample-complexities-for-solving-markov-decision-processes-with-a-generative-model
AUTHORS: Aaron Sidford, Mengdi Wang, Xian Wu, Lin Yang, Yinyu Ye
HIGHLIGHT: In this paper we consider the problem of computing an ϵ-optimal policy of a discounted Markov Decision Process (DMDP) provided we can only access its transition function through a generative sampling model that given any state-action pair samples from the transition function in $O(1)$ time.

480, TITLE: ChannelNets: Compact and Efficient Convolutional Neural Networks via Channel-Wise Convolutions
https://papers.nips.cc/paper/7766-channelnets-compact-and-efficient-convolutional-neural-networks-via-channel-wise-convolutions
AUTHORS: Hongyang Gao, Zhengyang Wang, Shuiwang Ji
HIGHLIGHT: In this work, we propose to compress deep models by using channel-wise convolutions, which replace dense connections among feature maps with sparse ones in CNNs.

481, TITLE: Causal Inference and Mechanism Clustering of A Mixture of Additive Noise Models
https://papers.nips.cc/paper/7767-causal-inference-and-mechanism-clustering-of-a-mixture-of-additive-noise-models
AUTHORS: Shoubo Hu, Zhitang Chen, Vahid Partovi Nia, Laiwan CHAN, Yanhui Geng
HIGHLIGHT: In this paper, we generalize the Additive Noise Model (ANM) to a mixture model, which consists of a finite number of ANMs, and provide the condition of its causal identifiability.

482, TITLE: Contour location via entropy reduction leveraging multiple information sources
https://papers.nips.cc/paper/7768-contour-location-via-entropy-reduction-leveraging-multiple-information-sources
AUTHORS: Alexandre Marques, Remi Lam, Karen Willcox
HIGHLIGHT: We introduce an algorithm to locate contours of functions that are expensive to evaluate.

483, TITLE: Assessing Generative Models via Precision and Recall
https://papers.nips.cc/paper/7769-assessing-generative-models-via-precision-and-recall
AUTHORS: Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, Sylvain Gelly
HIGHLIGHT: Recent advances in generative modeling have led to an increased interest in the study of statistical divergences as means of model comparison.

484, TITLE: Multiple-Step Greedy Policies in Approximate and Online Reinforcement Learning
https://papers.nips.cc/paper/7770-multiple-step-greedy-policies-in-approximate-and-online-reinforcement-learning
AUTHORS: Yonathan Efroni, Gal Dalal, Bruno Scherrer, Shie Mannor
HIGHLIGHT: In this work, we study multiple-step greedy algorithms in more practical setups.

485, TITLE: A Convex Duality Framework for GANs
https://papers.nips.cc/paper/7771-a-convex-duality-framework-for-gans
AUTHORS: Farzan Farnia, David Tse
HIGHLIGHT: In this work, we address this question by developing a convex duality framework for analyzing GANs.

486, TITLE: Horizon-Independent Minimax Linear Regression
https://papers.nips.cc/paper/7772-horizon-independent-minimax-linear-regression
AUTHORS: Alan Malek, Peter L. Bartlett
HIGHLIGHT: We consider online linear regression: at each round, an adversary reveals a covariate vector, the learner predicts a real value, the adversary reveals a label, and the learner suffers the squared prediction error.

487, TITLE: Exploiting Numerical Sparsity for Efficient Learning : Faster Eigenvector Computation and Regression
https://papers.nips.cc/paper/7773-exploiting-numerical-sparsity-for-efficient-learning-faster-eigenvector-computation-and-regression
AUTHORS: Neha Gupta, Aaron Sidford
HIGHLIGHT: In this paper, we obtain improved running times for regression and top eigenvector computation for numerically sparse matrices.

488, TITLE: Experimental Design for Cost-Aware Learning of Causal Graphs
https://papers.nips.cc/paper/7774-experimental-design-for-cost-aware-learning-of-causal-graphs
AUTHORS: Erik Lindgren, Murat Kocaoglu, Alexandros G. Dimakis, Sriram Vishwanath
HIGHLIGHT: We consider the minimum cost intervention design problem: Given the essential graph of a causal graph and a cost to intervene on a variable, identify the set of interventions with minimum total cost that can learn any causal graph with the given essential graph.

489, TITLE: Task-Driven Convolutional Recurrent Models of the Visual System
https://papers.nips.cc/paper/7775-task-driven-convolutional-recurrent-models-of-the-visual-system
AUTHORS: Aran Nayebi, Daniel Bear, Jonas Kubilius, Kohitij Kar, Surya Ganguli, David Sussillo, James J. DiCarlo, Daniel L. Yamins
HIGHLIGHT: Here we explored the role of recurrence in improving classification performance.

490, TITLE: Meta-Reinforcement Learning of Structured Exploration Strategies
https://papers.nips.cc/paper/7776-meta-reinforcement-learning-of-structured-exploration-strategies
AUTHORS: Abhishek Gupta, Russell Mendonca, YuXuan Liu, Pieter Abbeel, Sergey Levine
HIGHLIGHT: In this work, we study how prior tasks can inform an agent about how to explore effectively in new situations.

491, TITLE: Sample Efficient Stochastic Gradient Iterative Hard Thresholding Method for Stochastic Sparse Linear Regression with Limited Attribute Observation
https://papers.nips.cc/paper/7777-sample-efficient-stochastic-gradient-iterative-hard-thresholding-method-for-stochastic-sparse-linear-regression-with-limited-attribute-observation
AUTHORS: Tomoya Murata, Taiji Suzuki
HIGHLIGHT: We develop new stochastic gradient methods for efficiently solving sparse linear regression in a partial attribute observation setting, where learners are only allowed to observe a fixed number of actively chosen attributes per example at training and prediction times.

492, TITLE: Semi-supervised Deep Kernel Learning: Regression with Unlabeled Data by Minimizing Predictive Variance
https://papers.nips.cc/paper/7778-semi-supervised-deep-kernel-learning-regression-with-unlabeled-data-by-minimizing-predictive-variance
AUTHORS: Neal Jean, Sang Michael Xie, Stefano Ermon
HIGHLIGHT: We present semi-supervised deep kernel learning (SSDKL), a semi-supervised regression model based on minimizing predictive variance in the posterior regularization framework.

493, TITLE: Generalizing to Unseen Domains via Adversarial Data Augmentation
https://papers.nips.cc/paper/7779-generalizing-to-unseen-domains-via-adversarial-data-augmentation
AUTHORS: Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C. Duchi, Vittorio Murino, Silvio Savarese

HIGHLIGHT: Only using training data from a single source distribution, we propose an iterative procedure that augments the dataset with examples from a fictitious target domain that is "hard" under the current model.

494, TITLE: Hyperbolic Neural Networks
https://papers.nips.cc/paper/7780-hyperbolic-neural-networks
AUTHORS: Octavian Ganea, Gary Becigneul, Thomas Hofmann
HIGHLIGHT: Here, we bridge this gap in a principled manner by combining the formalism of Möbius gyrovector spaces with the Riemannian geometry of the Poincaré model of hyperbolic spaces.

495, TITLE: Breaking the Curse of Horizon: Infinite-Horizon Off-Policy Estimation
https://papers.nips.cc/paper/7781-breaking-the-curse-of-horizon-infinite-horizon-off-policy-estimation
AUTHORS: Qiang Liu, Lihong Li, Ziyang Tang, Dengyong Zhou
HIGHLIGHT: In this paper, we propose a new off-policy estimation method that applies IS directly on the stationary state-visitation distributions to avoid the exploding variance issue faced by existing estimators.Our key contribution is a novel approach to estimating the density ratio of two stationary distributions, with trajectories sampled from only the behavior distribution.

496, TITLE: Learning Task Specifications from Demonstrations
https://papers.nips.cc/paper/7782-learning-task-specifications-from-demonstrations
AUTHORS: Marcell Vazquez-Chanlatte, Susmit Jha, Ashish Tiwari, Mark K. Ho, Sanjit Seshia
HIGHLIGHT: In this paper, we formulate the specification inference task as a maximum a posteriori (MAP) probability inference problem, apply the principle of maximum entropy to derive an analytic demonstration likelihood model and give an efficient approach to search for the most likely specification in a large candidate pool of specifications.

497, TITLE: Learning a latent manifold of odor representations from neural responses in piriform cortex
https://papers.nips.cc/paper/7783-learning-a-latent-manifold-of-odor-representations-from-neural-responses-in-piriform-cortex
AUTHORS: Anqi Wu, Stan Pashkovski, Sandeep R. Datta, Jonathan W. Pillow
HIGHLIGHT: Here we use odor-evoked responses in piriform cortex to identify a latent manifold specifying latent distance relationships between olfactory stimuli.

498, TITLE: Fully Understanding The Hashing Trick
https://papers.nips.cc/paper/7784-fully-understanding-the-hashing-trick
AUTHORS: Casper B. Freksen, Lior Kamma, Kasper Green Larsen
HIGHLIGHT: Feature hashing, also known as {\em the hashing trick}, introduced by Weinberger et al. (2009), is one of the key techniques used in scaling-up machine learning algorithms.

499, TITLE: Evolved Policy Gradients
https://papers.nips.cc/paper/7785-evolved-policy-gradients
AUTHORS: Rein Houthooft, Yuhua Chen, Phillip Isola, Bradly Stadie, Filip Wolski, OpenAI Jonathan Ho, Pieter Abbeel
HIGHLIGHT: We propose a metalearning approach for learning gradient-based reinforcement learning (RL) algorithms.

500, TITLE: The Spectrum of the Fisher Information Matrix of a Single-Hidden-Layer Neural Network
https://papers.nips.cc/paper/7786-the-spectrum-of-the-fisher-information-matrix-of-a-single-hidden-layer-neural-network
AUTHORS: Jeffrey Pennington, Pratik Worah
HIGHLIGHT: In this work, we extend a recently-developed framework for studying spectra of nonlinear random matrices to characterize an important measure of curvature, namely the eigenvalues of the Fisher information matrix.

501, TITLE: Learning Concave Conditional Likelihood Models for Improved Analysis of Tandem Mass Spectra
https://papers.nips.cc/paper/7787-learning-concave-conditional-likelihood-models-for-improved-analysis-of-tandem-mass-spectra
AUTHORS: John T. Halloran, David M. Rocke
HIGHLIGHT: In this work, we greatly expand the parameter learning capabilities of a dynamic Bayesian network (DBN) peptide-scoring algorithm, Didea, by deriving emission distributions for which its conditional log-likelihood scoring function remains concave.

502, TITLE: Differentially Private k-Means with Constant Multiplicative Error
https://papers.nips.cc/paper/7788-differentially-private-k-means-with-constant-multiplicative-error
AUTHORS: Uri Stemmer, Haim Kaplan
HIGHLIGHT: We present, for the first time, efficient private algorithms for the problem with constant multiplicative error.

503, TITLE: Policy Optimization via Importance Sampling
https://papers.nips.cc/paper/7789-policy-optimization-via-importance-sampling

AUTHORS: Alberto Maria Metelli, Matteo Papini, Francesco Faccio, Marcello Restelli
HIGHLIGHT: In this paper, we propose a novel, model-free, policy search algorithm, POIS, applicable in both action-based and parameter-based settings.

504, TITLE: Estimating Learnability in the Sublinear Data Regime
https://papers.nips.cc/paper/7790-estimating-learnability-in-the-sublinear-data-regime
AUTHORS: Weihao Kong, Gregory Valiant
HIGHLIGHT: We consider the problem of estimating how well a model class is capable of fitting a distribution of labeled data.

505, TITLE: Algorithmic Assurance: An Active Approach to Algorithmic Testing using Bayesian Optimisation
https://papers.nips.cc/paper/7791-algorithmic-assurance-an-active-approach-to-algorithmic-testing-using-bayesian-optimisation
AUTHORS: Shivapratap Gopakumar, Sunil Gupta, Santu Rana, Vu Nguyen, Svetha Venkatesh
HIGHLIGHT: We introduce algorithmic assurance, the problem of testing whether machine learning algorithms are conforming to their intended design goal.

506, TITLE: Community Exploration: From Offline Optimization to Online Learning
https://papers.nips.cc/paper/7792-community-exploration-from-offline-optimization-to-online-learning
AUTHORS: Xiaowei Chen, Weiran Huang, Wei Chen, John C. S. Lui
HIGHLIGHT: We introduce the community exploration problem that has various real-world applications such as online advertising.

507, TITLE: A Dual Framework for Low-rank Tensor Completion
https://papers.nips.cc/paper/7793-a-dual-framework-for-low-rank-tensor-completion
AUTHORS: Madhav Nimishakavi, Pratik Kumar Jawanpuria, Bamdev Mishra
HIGHLIGHT: In this work, we fill this gap by proposing a variant of the latent trace norm that helps in learning a non-sparse combination of tensors.

508, TITLE: Low-rank Interaction with Sparse Additive Effects Model for Large Data Frames
https://papers.nips.cc/paper/7794-low-rank-interaction-with-sparse-additive-effects-model-for-large-data-frames
AUTHORS: Genevi?ve Robin, Hoi-To Wai, Julie Josse, Olga Klopp, Eric Moulines
HIGHLIGHT: In this paper, we introduce a low-rank interaction and sparse additive effects (LORIS) model which combines matrix regression on a dictionary and low-rank design, to estimate main effects and interactions simultaneously.

509, TITLE: Inference Aided Reinforcement Learning for Incentive Mechanism Design in Crowdsourcing
https://papers.nips.cc/paper/7795-inference-aided-reinforcement-learning-for-incentive-mechanism-design-in-crowdsourcing
AUTHORS: Zehong Hu, Yitao Liang, Jie Zhang, Zhao Li, Yang Liu
HIGHLIGHT: In this paper, we propose a novel inference aided reinforcement mechanism that acquires data sequentially and requires no such prior assumptions.

510, TITLE: Middle-Out Decoding
https://papers.nips.cc/paper/7796-middle-out-decoding
AUTHORS: Shikib Mehri, Leonid Sigal
HIGHLIGHT: In this paper, we speculate that a fundamental shortcoming of sequence generation models is that the decoding is done strictly from left-to-right, meaning that outputs values generated earlier have a profound effect on those generated later.

511, TITLE: First-order Stochastic Algorithms for Escaping From Saddle Points in Almost Linear Time
https://papers.nips.cc/paper/7797-first-order-stochastic-algorithms-for-escaping-from-saddle-points-in-almost-linear-time
AUTHORS: Yi Xu, Jing Rong, Tianbao Yang
HIGHLIGHT: (This is a theory paper) In this paper, we consider first-order methods for solving stochastic non-convex optimization problems.

512, TITLE: To Trust Or Not To Trust A Classifier
https://papers.nips.cc/paper/7798-to-trust-or-not-to-trust-a-classifier
AUTHORS: Heinrich Jiang, Been Kim, Melody Guan, Maya Gupta
HIGHLIGHT: We propose a new score, called the {\it trust score}, which measures the agreement between the classifier and a modified nearest-neighbor classifier on the testing example.

513, TITLE: Reparameterization Gradient for Non-differentiable Models
https://papers.nips.cc/paper/7799-reparameterization-gradient-for-non-differentiable-models

AUTHORS: Wonyeol Lee, Hangyeol Yu, Hongseok Yang
HIGHLIGHT: We present a new algorithm for stochastic variational inference that targets at models with non-differentiable densities.

514, TITLE: A Simple Proximal Stochastic Gradient Method for Nonsmooth Nonconvex Optimization
https://papers.nips.cc/paper/7800-a-simple-proximal-stochastic-gradient-method-for-nonsmooth-nonconvex-optimization
AUTHORS: Zhize Li, Jian Li
HIGHLIGHT: We propose a proximal stochastic gradient algorithm based on variance reduction, called ProxSVRG+.

515, TITLE: Multimodal Generative Models for Scalable Weakly-Supervised Learning
https://papers.nips.cc/paper/7801-multimodal-generative-models-for-scalable-weakly-supervised-learning
AUTHORS: Mike Wu, Noah Goodman
HIGHLIGHT: Here, we introduce a multimodal variational autoencoder (MVAE) that uses a product-of-experts inference network and a sub-sampled training paradigm to solve the multi-modal inference problem.

516, TITLE: How Much Restricted Isometry is Needed In Nonconvex Matrix Recovery?
https://papers.nips.cc/paper/7802-how-much-restricted-isometry-is-needed-in-nonconvex-matrix-recovery
AUTHORS: Richard Zhang, Cedric Josz, Somayeh Sojoudi, Javad Lavaei
HIGHLIGHT: In this paper, we show that moderate RIP is not enough to eliminate spurious local minima, so existing results can only hold for near-perfect RIP.

517, TITLE: Occam's razor is insufficient to infer the preferences of irrational agents
https://papers.nips.cc/paper/7803-occams-razor-is-insufficient-to-infer-the-preferences-of-irrational-agents
AUTHORS: Stuart Armstrong, S?ren Mindermann
HIGHLIGHT: Since human planning systematically deviates from rationality, several approaches have been tried to account for specific human shortcomings.

518, TITLE: Manifold Structured Prediction
https://papers.nips.cc/paper/7804-manifold-structured-prediction
AUTHORS: Alessandro Rudi, Carlo Ciliberto, GianMaria Marconi, Lorenzo Rosasco
HIGHLIGHT: Specifically, we study a structured prediction approach to manifold-valued regression.

519, TITLE: Fast Greedy MAP Inference for Determinantal Point Process to Improve Recommendation Diversity
https://papers.nips.cc/paper/7805-fast-greedy-map-inference-for-determinantal-point-process-to-improve-recommendation-diversity
AUTHORS: Laming Chen, Guoxin Zhang, Eric Zhou
HIGHLIGHT: To overcome the computational challenge, in this paper, we propose a novel algorithm to greatly accelerate the greedy MAP inference for DPP.

520, TITLE: Learning Others' Intentional Models in Multi-Agent Settings Using Interactive POMDPs
https://papers.nips.cc/paper/7806-learning-others-intentional-models-in-multi-agent-settings-using-interactive-pomdps
AUTHORS: Yanlin Han, Piotr Gmytrasiewicz
HIGHLIGHT: In order to predict other agents' actions using I-POMDPs, we propose an approach that effectively uses Bayesian inference and sequential Monte Carlo sampling to learn others' intentional models which ascribe to them beliefs, preferences and rationality in action selection.

521, TITLE: Contextual Pricing for Lipschitz Buyers
https://papers.nips.cc/paper/7807-contextual-pricing-for-lipschitz-buyers
AUTHORS: Jieming Mao, Renato Leme, Jon Schneider
HIGHLIGHT: We present improved bounds in the special case of a population of linear buyers.

522, TITLE: Online Improper Learning with an Approximation Oracle
https://papers.nips.cc/paper/7808-online-improper-learning-with-an-approximation-oracle
AUTHORS: Elad Hazan, Wei Hu, Yuanzhi Li, zhiyuan li
HIGHLIGHT: We study the following question: given an efficient approximation algorithm for an optimization problem, can we learn efficiently in the same setting?

523, TITLE: Bandit Learning in Concave N-Person Games
https://papers.nips.cc/paper/7809-bandit-learning-in-concave-n-person-games
AUTHORS: Mario Bravo, David Leslie, Panayotis Mertikopoulos
HIGHLIGHT: This paper examines the long-run behavior of learning with bandit feedback in non-cooperative concave games.

524, TITLE: On Fast Leverage Score Sampling and Optimal Learning
https://papers.nips.cc/paper/7810-on-fast-leverage-score-sampling-and-optimal-learning
AUTHORS: Alessandro Rudi, Daniele Calandriello, Luigi Carratino, Lorenzo Rosasco
HIGHLIGHT: In this paper, we study the problem of leverage score sampling for positive definite ma- trices defined by a kernel.

525, TITLE: Unsupervised Video Object Segmentation for Deep Reinforcement Learning
https://papers.nips.cc/paper/7811-unsupervised-video-object-segmentation-for-deep-reinforcement-learning
AUTHORS: Vikash Goel, Jameson Weng, Pascal Poupart
HIGHLIGHT: We present a new technique for deep reinforcement learning that automatically detects moving objects and uses the relevant information for action selection.

526, TITLE: Efficient inference for time-varying behavior during learning
https://papers.nips.cc/paper/7812-efficient-inference-for-time-varying-behavior-during-learning
AUTHORS: Nicholas A. Roy, Ji Hyun Bak, Athena Akrami, Carlos Brody, Jonathan W. Pillow
HIGHLIGHT: To overcome these limitations, we propose a dynamic psychophysical model that efficiently tracks trial-to-trial changes in behavior over the course of training.

527, TITLE: Learning convex polytopes with margin
https://papers.nips.cc/paper/7813-learning-convex-polytopes-with-margin
AUTHORS: Lee-Ad Gottlieb, Eran Kaufman, Aryeh Kontorovich, Gabriel Nivasch
HIGHLIGHT: We present improved algorithm for properly learning convex polytopes in the realizable PAC setting from data with a margin.

528, TITLE: Critical initialisation for deep signal propagation in noisy rectifier neural networks
https://papers.nips.cc/paper/7814-critical-initialisation-for-deep-signal-propagation-in-noisy-rectifier-neural-networks
AUTHORS: Arnu Pretorius, Elan van Biljon, Steve Kroon, Herman Kamper
HIGHLIGHT: We use this framework to investigate initialisation strategies for noisy ReLU networks.

529, TITLE: Insights on representational similarity in neural networks with canonical correlation
https://papers.nips.cc/paper/7815-insights-on-representational-similarity-in-neural-networks-with-canonical-correlation
AUTHORS: Ari Morcos, Maithra Raghu, Samy Bengio
HIGHLIGHT: Here, we develop projection weighted CCA (Canonical Correlation Analysis) as a tool for understanding neural networks, building off of SVCCA, a recently proposed method (Raghu et al, 2017).

530, TITLE: Variational Inference with Tail-adaptive f-Divergence
https://papers.nips.cc/paper/7816-variational-inference-with-tail-adaptive-f-divergence
AUTHORS: Dilin Wang, Hao Liu, Qiang Liu
HIGHLIGHT: In this paper, we propose a new class of tail-adaptive f-divergences that adaptively change the convex function f with the tail of the importance weights, in a way that theoretically guarantee finite moments, while simultaneously achieving mass-covering properties.

531, TITLE: Mental Sampling in Multimodal Representations
https://papers.nips.cc/paper/7817-mental-sampling-in-multimodal-representations
AUTHORS: Jianqiao Zhu, Adam Sanborn, Nick Chater
HIGHLIGHT: To describe people's internal and external foraging behavior, various random walk models have been proposed.

532, TITLE: Adversarially Robust Optimization with Gaussian Processes
https://papers.nips.cc/paper/7818-adversarially-robust-optimization-with-gaussian-processes
AUTHORS: Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, Volkan Cevher
HIGHLIGHT: In this paper, we consider the problem of Gaussian process (GP) optimization with an added robustness requirement: The returned point may be perturbed by an adversary, and we require the function value to remain as high as possible even after this perturbation.

533, TITLE: Learning to Multitask
https://papers.nips.cc/paper/7819-learning-to-multitask
AUTHORS: Yu Zhang, Ying Wei, Qiang Yang
HIGHLIGHT: In order to identify an effective multitask model for a given multitask problem, we propose a learning framework called Learning to MultiTask (L2MT).

534, TITLE: Loss Functions for Multiset Prediction
https://papers.nips.cc/paper/7820-loss-functions-for-multiset-prediction
AUTHORS: Sean Welleck, Zixin Yao, Yu Gai, Jialin Mao, Zheng Zhang, Kyunghyun Cho
HIGHLIGHT: In this paper, we propose a novel multiset loss function by viewing this problem from the perspective of sequential decision making.

535, TITLE: Computing Kantorovich-Wasserstein Distances on <var>d</var>-dimensional histograms using <var>(d+1)</var>-partite graphs
https://papers.nips.cc/paper/7821-computing-kantorovich-wasserstein-distances-on-d-dimensional-histograms-using-d1-partite-graphs
AUTHORS: Gennaro Auricchio, Federico Bassetti, Stefano Gualandi, Marco Veneroni
HIGHLIGHT: This paper presents a novel method to compute the exact Kantorovich-Wasserstein distance between a pair of d-dimensional histograms having n bins each.

536, TITLE: Neural Interaction Transparency (NIT): Disentangling Learned Interactions for Improved Interpretability
https://papers.nips.cc/paper/7822-neural-interaction-transparency-nit-disentangling-learned-interactions-for-improved-interpretability
AUTHORS: Michael Tsang, Hanpeng Liu, Sanjay Purushotham, Pavankumar Murali, Yan Liu
HIGHLIGHT: We propose a framework, Neural Interaction Transparency (NIT), that disentangles the shared learning across different interactions to obtain their intrinsic lower-order and interpretable structure.

537, TITLE: CapProNet: Deep Feature Learning via Orthogonal Projections onto Capsule Subspaces
https://papers.nips.cc/paper/7823-cappronet-deep-feature-learning-via-orthogonal-projections-onto-capsule-subspaces
AUTHORS: Liheng Zhang, Marzieh Edraki, Guo-Jun Qi
HIGHLIGHT: In this paper, we formalize the idea behind capsule nets of using a capsule vector rather than a neuron activation to predict the label of samples.

538, TITLE: Gamma-Poisson Dynamic Matrix Factorization Embedded with Metadata Influence
https://papers.nips.cc/paper/7824-gamma-poisson-dynamic-matrix-factorization-embedded-with-metadata-influence
AUTHORS: Trong Dinh Thac Do, Longbing Cao
HIGHLIGHT: A conjugate Gamma-Poisson model for Dynamic Matrix Factorization incorporated with metadata influence (mGDMF for short) is proposed to effectively and efficiently model massive, sparse and dynamic data in recommendations.

539, TITLE: Masking: A New Perspective of Noisy Supervision
https://papers.nips.cc/paper/7825-masking-a-new-perspective-of-noisy-supervision
AUTHORS: Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, Masashi Sugiyama
HIGHLIGHT: In this paper, we propose a human-assisted approach called "Masking" that conveys human cognition of invalid class transitions and naturally speculates the structure of the noise transition matrix.

540, TITLE: On GANs and GMMs
https://papers.nips.cc/paper/7826-on-gans-and-gmms
AUTHORS: Eitan Richardson, Yair Weiss
HIGHLIGHT: In this paper, we examine the utility of GANs in learning statistical models of images by comparing them to perhaps the simplest statistical model, the Gaussian Mixture Model.

541, TITLE: Differential Properties of Sinkhorn Approximation for Learning with Wasserstein Distance
https://papers.nips.cc/paper/7827-differential-properties-of-sinkhorn-approximation-for-learning-with-wasserstein-distance
AUTHORS: Giulia Luise, Alessandro Rudi, Massimiliano Pontil, Carlo Ciliberto
HIGHLIGHT: In this work we characterize the differential properties of the original Sinkhorn approximation, proving that it enjoys the same smoothness as its regularized version and we explicitly provide an efficient algorithm to compute its gradient.

542, TITLE: Practical Deep Stereo (PDS): Toward applications-friendly deep stereo matching
https://papers.nips.cc/paper/7828-practical-deep-stereo-pds-toward-applications-friendly-deep-stereo-matching
AUTHORS: Stepan Tulyakov, Anton Ivanov, Fran?ois Fleuret
HIGHLIGHT: The Practical Deep Stereo (PDS) network that we propose addresses both issues: First, its architecture relies on novel bottleneck modules that drastically reduce the memory footprint in inference, and additional design choices allow to handle greater image size during training.

543, TITLE: A Bayes-Sard Cubature Method
https://papers.nips.cc/paper/7829-a-bayes-sard-cubature-method
AUTHORS: Toni Karvonen, Chris J. Oates, Simo Sarkka

HIGHLIGHT: To address these drawbacks we introduce Bayes-Sard cubature, a probabilistic framework that combines the flexibility of Bayesian cubature with the robustness of classical cubatures which are well-established.

544, TITLE: Dual Swap Disentangling
https://papers.nips.cc/paper/7830-dual-swap-disentangling
AUTHORS: Zunlei Feng, Xinchao Wang, Chenglong Ke, An-Xiang Zeng, Dacheng Tao, Mingli Song
HIGHLIGHT: In this paper, we propose a weakly semi-supervised method, termed as Dual Swap Disentangling (DSD), for disentangling using both labeled and unlabeled data.

545, TITLE: Diverse Ensemble Evolution: Curriculum Data-Model Marriage
https://papers.nips.cc/paper/7831-diverse-ensemble-evolution-curriculum-data-model-marriage
AUTHORS: Tianyi Zhou, Shengjie Wang, Jeff A. Bilmes
HIGHLIGHT: We study a new method (``Diverse Ensemble Evolution (DivE^2)") to train an ensemble of machine learning models that assigns data to models at each training epoch based on each model's current expertise and an intra- and inter-model diversity reward.

546, TITLE: Binary Classification from Positive-Confidence Data
https://papers.nips.cc/paper/7832-binary-classification-from-positive-confidence-data
AUTHORS: Takashi Ishida, Gang Niu, Masashi Sugiyama
HIGHLIGHT: For the Pconf classification problem, we provide a simple empirical risk minimization framework that is model-independent and optimization-independent.

547, TITLE: Deep Generative Models for Distribution-Preserving Lossy Compression
https://papers.nips.cc/paper/7833-deep-generative-models-for-distribution-preserving-lossy-compression
AUTHORS: Michael Tschannen, Eirikur Agustsson, Mario Lucic
HIGHLIGHT: We propose and study the problem of distribution-preserving lossy compression.

548, TITLE: Exact natural gradient in deep linear networks and its application to the nonlinear case
https://papers.nips.cc/paper/7834-exact-natural-gradient-in-deep-linear-networks-and-its-application-to-the-nonlinear-case
AUTHORS: Alberto Bernacchia, Mate Lengyel, Guillaume Hennequin
HIGHLIGHT: We provide for the first time an analytical solution for its convergence rate, showing that the loss decreases exponentially to the global minimum in parameter space.

549, TITLE: Constructing Fast Network through Deconstruction of Convolution
https://papers.nips.cc/paper/7835-constructing-fast-network-through-deconstruction-of-convolution
AUTHORS: Yunho Jeon, Junmo Kim
HIGHLIGHT: To cope with various convolutions, we propose a new shift operation called active shift layer (ASL) that formulates the amount of shift as a learnable function with shift parameters.

550, TITLE: Memory Replay GANs: Learning to Generate New Categories without Forgetting
https://papers.nips.cc/paper/7836-memory-replay-gans-learning-to-generate-new-categories-without-forgetting
AUTHORS: Chenshen Wu, Luis Herranz, Xialei Liu, yaxing wang, Joost van de Weijer, Bogdan Raducanu
HIGHLIGHT: In this paper we consider the case of generative models.

551, TITLE: The Convergence of Sparsified Gradient Methods
https://papers.nips.cc/paper/7837-the-convergence-of-sparsified-gradient-methods
AUTHORS: Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, Cedric Renggli
HIGHLIGHT: This is the question we address in this paper.

552, TITLE: Automating Bayesian optimization with Bayesian optimization
https://papers.nips.cc/paper/7838-automating-bayesian-optimization-with-bayesian-optimization
AUTHORS: Gustavo Malkomes, Roman Garnett
HIGHLIGHT: In this work, we introduce a novel automated Bayesian optimization approach that dynamically selects promising models for explaining the observed data using Bayesian Optimization in the model space.

553, TITLE: Stacked Semantics-Guided Attention Model for Fine-Grained Zero-Shot Learning
https://papers.nips.cc/paper/7839-stacked-semantics-guided-attention-model-for-fine-grained-zero-shot-learning
AUTHORS: yunlong yu, Zhong Ji, Yanwei Fu, Jichang Guo, Yanwei Pang, Zhongfei (Mark) Zhang

HIGHLIGHT: To this end, we propose a novel stacked semantics-guided attention (S2GA) model to obtain semantic relevant features by using individual class semantic features to progressively guide the visual features to generate an attention map for weighting the importance of different local regions.

554, TITLE: Dirichlet-based Gaussian Processes for Large-scale Calibrated Classification
https://papers.nips.cc/paper/7840-dirichlet-based-gaussian-processes-for-large-scale-calibrated-classification
AUTHORS: Dimitrios Milios, Raffaello Camoriano, Pietro Michiardi, Lorenzo Rosasco, Maurizio Filippone
HIGHLIGHT: In this work, we investigate if and how Gaussian process regression directly applied to classification labels can be used to tackle this question.

555, TITLE: Multi-Task Zipping via Layer-wise Neuron Sharing
https://papers.nips.cc/paper/7841-multi-task-zipping-via-layer-wise-neuron-sharing
AUTHORS: Xiaoxi He, Zimu Zhou, Lothar Thiele
HIGHLIGHT: In this work, we aim to reduce the redundancy across multiple models.

556, TITLE: Dimensionally Tight Bounds for Second-Order Hamiltonian Monte Carlo
https://papers.nips.cc/paper/7842-dimensionally-tight-bounds-for-second-order-hamiltonian-monte-carlo
AUTHORS: Oren Mangoubi, Nisheeth Vishnoi
HIGHLIGHT: Here we show that this conjecture is true when sampling from strongly log-concave target distributions that satisfy a weak third-order regularity property associated with the input data.

557, TITLE: Approximation algorithms for stochastic clustering
https://papers.nips.cc/paper/7843-approximation-algorithms-for-stochastic-clustering
AUTHORS: David Harris, Shi Li, Aravind Srinivasan, Khoa Trinh, Thomas Pensyl
HIGHLIGHT: We consider stochastic settings for clustering, and develop provably-good (approximation) algorithms for a number of these notions.

558, TITLE: Evolutionary Stochastic Gradient Descent for Optimization of Deep Neural Networks
https://papers.nips.cc/paper/7844-evolutionary-stochastic-gradient-descent-for-optimization-of-deep-neural-networks
AUTHORS: Xiaodong Cui, Wei Zhang, Zolt?n T?ske, Michael Picheny
HIGHLIGHT: We propose a population-based Evolutionary Stochastic Gradient Descent (ESGD) framework for optimizing deep neural networks.

559, TITLE: Learning to Infer Graphics Programs from Hand-Drawn Images
https://papers.nips.cc/paper/7845-learning-to-infer-graphics-programs-from-hand-drawn-images
AUTHORS: Kevin Ellis, Daniel Ritchie, Armando Solar-Lezama, Josh Tenenbaum
HIGHLIGHT: We introduce a model that learns to convert simple hand drawings into graphics programs written in a subset of \LaTeX.

560, TITLE: Graphical Generative Adversarial Networks
https://papers.nips.cc/paper/7846-graphical-generative-adversarial-networks
AUTHORS: Chongxuan LI, Max Welling, Jun Zhu, Bo Zhang
HIGHLIGHT: We propose Graphical Generative Adversarial Networks (Graphical-GAN) to model structured data.

561, TITLE: Variational Learning on Aggregate Outputs with Gaussian Processes
https://papers.nips.cc/paper/7847-variational-learning-on-aggregate-outputs-with-gaussian-processes
AUTHORS: Ho Chung Law, Dino Sejdinovic, Ewan Cameron, Tim Lucas, Seth Flaxman, Katherine Battle, Kenji Fukumizu
HIGHLIGHT: We propose new bounds and tractable approximations, leading to improved prediction accuracy and scalability to large datasets, while explicitly taking uncertainty into account.

562, TITLE: MacNet: Transferring Knowledge from Machine Comprehension to Sequence-to-Sequence Models
https://papers.nips.cc/paper/7848-macnet-transferring-knowledge-from-machine-comprehension-to-sequence-to-sequence-models
AUTHORS: Boyuan Pan, Yazheng Yang, Hao Li, Zhou Zhao, Yueting Zhuang, Deng Cai, Xiaofei He
HIGHLIGHT: In this paper, we transfer knowledge learned from machine comprehension to the sequence-to-sequence tasks to deepen the understanding of the text.

563, TITLE: Poison Frogs! Targeted Clean-Label Poisoning Attacks on Neural Networks
https://papers.nips.cc/paper/7849-poison-frogs-targeted-clean-label-poisoning-attacks-on-neural-networks
AUTHORS: Ali Shafahi, W. Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, Tom Goldstein

HIGHLIGHT: We present an optimization-based method for crafting poisons, and show that just one single poison image can control classifier behavior when transfer learning is used.

564, TITLE: Information Constraints on Auto-Encoding Variational Bayes
https://papers.nips.cc/paper/7850-information-constraints-on-auto-encoding-variational-bayes
AUTHORS: Romain Lopez, Jeffrey Regier, Michael I. Jordan, Nir Yosef
HIGHLIGHT: We propose a framework for learning representations that relies on Auto-Encoding Variational Bayes and whose search space is constrained via kernel-based measures of independence.

565, TITLE: Recurrent Transformer Networks for Semantic Correspondence
https://papers.nips.cc/paper/7851-recurrent-transformer-networks-for-semantic-correspondence
AUTHORS: Seungryong Kim, Stephen Lin, SANG RYUL JEON, Dongbo Min, Kwanghoon Sohn
HIGHLIGHT: We present recurrent transformer networks (RTNs) for obtaining dense correspondences between semantically similar images.

566, TITLE: Online convex optimization for cumulative constraints
https://papers.nips.cc/paper/7852-online-convex-optimization-for-cumulative-constraints
AUTHORS: Jianjun Yuan, Andrew Lamperski
HIGHLIGHT: We propose the algorithms for online convex optimization which lead to cumulative squared constraint violations of the form $\sum\limits_{t=1}^T\big([g(x_t)]_+\big)^2=O(T^{1-\beta})$, where $\beta\in(0,1)$.

567, TITLE: Predict Responsibly: Improving Fairness and Accuracy by Learning to Defer
https://papers.nips.cc/paper/7853-predict-responsibly-improving-fairness-and-accuracy-by-learning-to-defer
AUTHORS: David Madras, Toni Pitassi, Richard Zemel
HIGHLIGHT: In this work, we explore a simple version of this interaction with a two-stage framework containing an automated model and an external decision-maker.

568, TITLE: Deep State Space Models for Unconditional Word Generation
https://papers.nips.cc/paper/7854-deep-state-space-models-for-unconditional-word-generation
AUTHORS: Florian Schmidt, Thomas Hofmann
HIGHLIGHT: Autoregressive feedback is considered a necessity for successful unconditional text generation using stochastic sequence models.

569, TITLE: ResNet with one-neuron hidden layers is a Universal Approximator
https://papers.nips.cc/paper/7855-resnet-with-one-neuron-hidden-layers-is-a-universal-approximator
AUTHORS: Hongzhou Lin, Stefanie Jegelka
HIGHLIGHT: We demonstrate that a very deep ResNet with stacked modules that have one neuron per hidden layer and ReLU activation functions can uniformly approximate any Lebesgue integrable function in d dimensions, i.e. \ell_1(R^d).

570, TITLE: Transfer of Value Functions via Variational Methods
https://papers.nips.cc/paper/7856-transfer-of-value-functions-via-variational-methods
AUTHORS: Andrea Tirinzoni, Rafael Rodriguez Sanchez, Marcello Restelli
HIGHLIGHT: We propose an approach that uses the given source tasks to learn a prior distribution over optimal value functions and provide an efficient variational approximation of the corresponding posterior in a new target task.

571, TITLE: The Cluster Description Problem - Complexity Results, Formulations and Approximations
https://papers.nips.cc/paper/7857-the-cluster-description-problem-complexity-results-formulations-and-approximations
AUTHORS: Ian Davidson, Antoine Gourru, S Ravi
HIGHLIGHT: A challenge for explainable AI is to find a compact and distinct explanations of each cluster which in this paper is using instance-level descriptors/tags from a common dictionary.

572, TITLE: Sharp Bounds for Generalized Uniformity Testing
https://papers.nips.cc/paper/7858-sharp-bounds-for-generalized-uniformity-testing
AUTHORS: Ilias Diakonikolas, Daniel M. Kane, Alistair Stewart
HIGHLIGHT: In more detail, we present a computationally efficient tester whose sample complexity is optimal, within constant factors, and a matching worst-case information-theoretic lower bound.

573, TITLE: Deep Neural Networks with Box Convolutions
https://papers.nips.cc/paper/7859-deep-neural-networks-with-box-convolutions
AUTHORS: Egor Burkov, Victor Lempitsky

HIGHLIGHT: Here, we show that a convolutional layer that computes box filter responses in a sliding manner can be used within deep architectures, whereas the dimensions and the offsets of the sliding boxes in such a layer can be learned as part of an end-to-end loss minimization.

574, TITLE: Learning towards Minimum Hyperspherical Energy
https://papers.nips.cc/paper/7860-learning-towards-minimum-hyperspherical-energy
AUTHORS: Weiyang Liu, Rongmei Lin, Zhen Liu, Lixin Liu, Zhiding Yu, Bo Dai, Le Song
HIGHLIGHT: In light of this intuition, we reduce the redundancy regularization problem to generic energy minimization, and propose a minimum hyperspherical energy (MHE) objective as generic regularization for neural networks.

575, TITLE: LF-Net: Learning Local Features from Images
https://papers.nips.cc/paper/7861-lf-net-learning-local-features-from-images
AUTHORS: Yuki Ono, Eduard Trulls, Pascal Fua, Kwang Moo Yi
HIGHLIGHT: We present a novel deep architecture and a training strategy to learn a local feature pipeline from scratch, using collections of images without the need for human supervision.

576, TITLE: SLANG: Fast Structured Covariance Approximations for Bayesian Deep Learning with Natural Gradient
https://papers.nips.cc/paper/7862-slang-fast-structured-covariance-approximations-for-bayesian-deep-learning-with-natural-gradient
AUTHORS: Aaron Mishkin, Frederik Kunstner, Didrik Nielsen, Mark Schmidt, Mohammad Emtiyaz Khan
HIGHLIGHT: To address this issue, we propose a new stochastic, low-rank, approximate natural-gradient (SLANG) method for variational inference in large deep models.

577, TITLE: Tangent: Automatic differentiation using source-code transformation for dynamically typed array programming
https://papers.nips.cc/paper/7863-tangent-automatic-differentiation-using-source-code-transformation-for-dynamically-typed-array-programming
AUTHORS: Bart van Merrienboer, Dan Moldovan, Alexander Wiltschko
HIGHLIGHT: In this work, we explore techniques from the field of automatic differentiation (AD) that can give researchers expressive power, performance and strong usability.

578, TITLE: Multi-domain Causal Structure Learning in Linear Systems
https://papers.nips.cc/paper/7864-multi-domain-causal-structure-learning-in-linear-systems
AUTHORS: AmirEmad Ghassami, Negar Kiyavash, Biwei Huang, Kun Zhang
HIGHLIGHT: Our approach unifies the idea in those works and generalizes to the case that there is no such invariance across the domains.

579, TITLE: Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences
https://papers.nips.cc/paper/7865-privacy-amplification-by-subsampling-tight-analyses-via-couplings-and-divergences
AUTHORS: Borja Balle, Gilles Barthe, Marco Gaboardi
HIGHLIGHT: In this paper we present a general method that recovers and improves prior analyses, yields lower bounds and derives new instances of privacy amplification by subsampling.

580, TITLE: Exponentially Weighted Imitation Learning for Batched Historical Data
https://papers.nips.cc/paper/7866-exponentially-weighted-imitation-learning-for-batched-historical-data
AUTHORS: Qing Wang, Jiechao Xiong, Lei Han, peng sun, Han Liu, Tong Zhang
HIGHLIGHT: To solve this problem, we propose a monotonic advantage reweighted imitation learning strategy that is applicable to problems with complex nonlinear function approximation and works well with hybrid (discrete and continuous) action space.

581, TITLE: Algebraic tests of general Gaussian latent tree models
https://papers.nips.cc/paper/7867-algebraic-tests-of-general-gaussian-latent-tree-models
AUTHORS: Dennis Leung, Mathias Drton
HIGHLIGHT: Illustrating with the star tree, we propose a new testing methodology that circumvents singularity issues by trading off some statistical estimation efficiency and handles cases with many constraints through recent advances on Gaussian approximation for maxima of sums of high-dimensional random vectors.

582, TITLE: Navigating with Graph Representations for Fast and Scalable Decoding of Neural Language Models
https://papers.nips.cc/paper/7868-navigating-with-graph-representations-for-fast-and-scalable-decoding-of-neural-language-models
AUTHORS: Minjia Zhang, Wenhan Wang, Xiaodong Liu, Jianfeng Gao, Yuxiong He
HIGHLIGHT: This paper proposes a novel softmax layer approximation algorithm, called Fast Graph Decoder (FGD), which quickly identifies, for a given context, a set of K words that are most likely to occur according to a NLM.

583, TITLE: Deep Structured Prediction with Nonlinear Output Transformations
https://papers.nips.cc/paper/7869-deep-structured-prediction-with-nonlinear-output-transformations
AUTHORS: Colin Graber, Ofer Meshi, Alexander Schwing
HIGHLIGHT: Here, we develop a novel model which generalizes existing approaches, such as structured prediction energy networks, and discuss a formulation which maintains applicability of existing inference techniques.

584, TITLE: Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling
https://papers.nips.cc/paper/7870-sequential-test-for-the-lowest-mean-from-thompson-to-murphy-sampling
AUTHORS: Emilie Kaufmann, Wouter M. Koolen, Aur?lien Garivier
HIGHLIGHT: We formalize this learning task as the problem of sequentially testing how the minimum mean among a finite set of distributions compares to a given threshold.

585, TITLE: Distributed Learning without Distress: Privacy-Preserving Empirical Risk Minimization
https://papers.nips.cc/paper/7871-distributed-learning-without-distress-privacy-preserving-empirical-risk-minimization
AUTHORS: Bargav Jayaraman, Lingxiao Wang, David Evans, Quanquan Gu
HIGHLIGHT: We present a distributed learning approach that combines differential privacy with secure multi-party computation.

586, TITLE: A no-regret generalization of hierarchical softmax to extreme multi-label classification
https://papers.nips.cc/paper/7872-a-no-regret-generalization-of-hierarchical-softmax-to-extreme-multi-label-classification
AUTHORS: Marek Wydmuch, Kalina Jasinska, Mikhail Kuznetsov, R?bert Busa-Fekete, Krzysztof Dembczynski
HIGHLIGHT: In this paper, we investigate probabilistic label trees (PLTs) that have been recently devised for tackling XMLC problems.

587, TITLE: Efficient Formal Safety Analysis of Neural Networks
https://papers.nips.cc/paper/7873-efficient-formal-safety-analysis-of-neural-networks
AUTHORS: Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, Suman Jana
HIGHLIGHT: In this paper, we present a new efficient approach for rigorously checking different safety properties of neural networks that significantly outperforms existing approaches by multiple orders of magnitude.

588, TITLE: Bayesian Distributed Stochastic Gradient Descent
https://papers.nips.cc/paper/7874-bayesian-distributed-stochastic-gradient-descent
AUTHORS: Michael Teng, Frank Wood
HIGHLIGHT: We introduce Bayesian distributed stochastic gradient descent (BDSGD), a high-throughput algorithm for training deep neural networks on parallel clusters.

589, TITLE: Visualizing the Loss Landscape of Neural Nets
https://papers.nips.cc/paper/7875-visualizing-the-loss-landscape-of-neural-nets
AUTHORS: Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, Tom Goldstein
HIGHLIGHT: In this paper, we explore the structure of neural loss functions, and the effect of loss landscapes on generalization, using a range of visualization methods.

590, TITLE: The Limits of Post-Selection Generalization
https://papers.nips.cc/paper/7876-the-limits-of-post-selection-generalization
AUTHORS: Jonathan Ullman, Adam Smith, Kobbi Nissim, Uri Stemmer, Thomas Steinke
HIGHLIGHT: In this work we show several limitations on the power of algorithms satisfying post hoc generalization.

591, TITLE: Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation
https://papers.nips.cc/paper/7877-graph-convolutional-policy-network-for-goal-directed-molecular-graph-generation
AUTHORS: Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, Jure Leskovec
HIGHLIGHT: Here we propose Graph Convolutional Policy Network (GCPN), a general graph convolutional network based model for goal-directed graph generation through reinforcement learning.

592, TITLE: On Controllable Sparse Alternatives to Softmax
https://papers.nips.cc/paper/7878-on-controllable-sparse-alternatives-to-softmax
AUTHORS: Anirban Laha, Saneem Ahmed Chemmengath, Priyanka Agrawal, Mitesh Khapra, Karthik Sankaranarayanan, Harish G. Ramaswamy
HIGHLIGHT: Within this framework, we propose two novel sparse formulations, sparsegen-lin and sparsehourglass, that seek to provide a control over the degree of desired sparsity.

593, TITLE: L4: Practical loss-based stepsize adaptation for deep learning
https://papers.nips.cc/paper/7879-l4-practical-loss-based-stepsize-adaptation-for-deep-learning
AUTHORS: Michal Rolinek, Georg Martius
HIGHLIGHT: We propose a stepsize adaptation scheme for stochastic gradient descent.

594, TITLE: Learning Latent Subspaces in Variational Autoencoders
https://papers.nips.cc/paper/7880-learning-latent-subspaces-in-variational-autoencoders
AUTHORS: Jack Klys, Jake Snell, Richard Zemel
HIGHLIGHT: We propose a VAE-based generative model which we show is capable of extracting features correlated to binary labels in the data and structuring it in a latent subspace which is easy to interpret.

595, TITLE: Turbo Learning for CaptionBot and DrawingBot
https://papers.nips.cc/paper/7881-turbo-learning-for-captionbot-and-drawingbot
AUTHORS: Qiuyuan Huang, Pengchuan Zhang, Dapeng Wu, Lei Zhang
HIGHLIGHT: We study in this paper the problems of both image captioning and text-to-image generation, and present a novel turbo learning approach to jointly training an image-to-text generator (a.k.a. CaptionBot) and a text-to-image generator (a.k.a. DrawingBot).

596, TITLE: Learning to Teach with Dynamic Loss Functions
https://papers.nips.cc/paper/7882-learning-to-teach-with-dynamic-loss-functions
AUTHORS: Lijun Wu, Fei Tian, Yingce Xia, Yang Fan, Tao Qin, Lai Jian-Huang, Tie-Yan Liu
HIGHLIGHT: In this work, we explore the possibility of imitating human teaching behaviors by dynamically and automatically outputting appropriate loss functions to train machine learning models.

597, TITLE: Multi-View Silhouette and Depth Decomposition for High Resolution 3D Object Representation
https://papers.nips.cc/paper/7883-multi-view-silhouette-and-depth-decomposition-for-high-resolution-3d-object-representation
AUTHORS: Edward Smith, Scott Fujimoto, David Meger
HIGHLIGHT: Drawing motivation from the canonical view representation of objects, we introduce a novel method for the fast up-sampling of 3D objects in voxel space through networks that perform super-resolution on the six orthographic depth projections.

598, TITLE: Size-Noise Tradeoffs in Generative Networks
https://papers.nips.cc/paper/7884-size-noise-tradeoffs-in-generative-networks
AUTHORS: Bolton Bailey, Matus J. Telgarsky
HIGHLIGHT: This paper investigates the ability of generative networks to convert their input noise distributions into other distributions.

599, TITLE: Online Adaptive Methods, Universality and Acceleration
https://papers.nips.cc/paper/7885-online-adaptive-methods-universality-and-acceleration
AUTHORS: Yehuda Kfir Levy, Alp Yurtsever, Volkan Cevher
HIGHLIGHT: We present a novel method for convex unconstrained optimization that, without any modifications ensures: (1) accelerated convergence rate for smooth objectives, (2) standard convergence rate in the general (non-smooth) setting, and (3) standard convergence rate in the stochastic optimization setting.

600, TITLE: Compact Generalized Non-local Network
https://papers.nips.cc/paper/7886-compact-generalized-non-local-network
AUTHORS: Kaiyu Yue, Ming Sun, Yuchen Yuan, Feng Zhou, Errui Ding, Fuxin Xu
HIGHLIGHT: To address this limitation, we generalize the non-local module and take the correlations between the positions of any two channels into account.

601, TITLE: On the Local Hessian in Back-propagation
https://papers.nips.cc/paper/7887-on-the-local-hessian-in-back-propagation
AUTHORS: Huishuai Zhang, Wei Chen, Tie-Yan Liu
HIGHLIGHT: In this paper, we take one step towards clarifying this problem.

602, TITLE: The Everlasting Database: Statistical Validity at a Fair Price
https://papers.nips.cc/paper/7888-the-everlasting-database-statistical-validity-at-a-fair-price
AUTHORS: Blake E. Woodworth, Vitaly Feldman, Saharon Rosset, Nati Srebro
HIGHLIGHT: We propose a mechanism for answering an arbitrarily long sequence of potentially adaptive statistical queries, by charging a price for each query and using the proceeds to collect additional samples.

603, TITLE: Lipschitz-Margin Training: Scalable Certification of Perturbation Invariance for Deep Neural Networks
https://papers.nips.cc/paper/7889-lipschitz-margin-training-scalable-certification-of-perturbation-invariance-for-deep-neural-networks
AUTHORS: Yusuke Tsuzuku, Issei Sato, Masashi Sugiyama
HIGHLIGHT: From the relationship between the Lipschitz constants and prediction margins, we present a computationally efficient calculation technique to lower-bound the size of adversarial perturbations that can deceive networks, and that is widely applicable to various complicated networks.

604, TITLE: Proximal SCOPE for Distributed Sparse Learning
https://papers.nips.cc/paper/7890-proximal-scope-for-distributed-sparse-learning
AUTHORS: Shenyi Zhao, Gong-Duo Zhang, Ming-Wei Li, Wu-Jun Li
HIGHLIGHT: In this paper, we propose a novel method, called proximal SCOPE (pSCOPE), for distributed sparse learning with L1 regularization.

605, TITLE: On Coresets for Logistic Regression
https://papers.nips.cc/paper/7891-on-coresets-for-logistic-regression
AUTHORS: Alexander Munteanu, Chris Schwiegelshohn, Christian Sohler, David Woodruff
HIGHLIGHT: To deal with intractable worst-case instances we introduce a complexity measure $\mu(X)$, which quantifies the hardness of compressing a data set for logistic regression.

606, TITLE: Neural Ordinary Differential Equations
https://papers.nips.cc/paper/7892-neural-ordinary-differential-equations
AUTHORS: Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, David K. Duvenaud
HIGHLIGHT: We introduce a new family of deep neural network models.

607, TITLE: Unsupervised Learning of Artistic Styles with Archetypal Style Analysis
https://papers.nips.cc/paper/7893-unsupervised-learning-of-artistic-styles-with-archetypal-style-analysis
AUTHORS: Daan Wynen, Cordelia Schmid, Julien Mairal
HIGHLIGHT: In this paper, we introduce an unsupervised learning approach to automatically dis- cover, summarize, and manipulate artistic styles from large collections of paintings.

608, TITLE: Approximating Real-Time Recurrent Learning with Random Kronecker Factors
https://papers.nips.cc/paper/7894-approximating-real-time-recurrent-learning-with-random-kronecker-factors
AUTHORS: Asier Mujika, Florian Meier, Angelika Steger
HIGHLIGHT: In this paper we propose the Kronecker Factored RTRL (KF-RTRL) algorithm that uses a Kronecker product decomposition to approximate the gradients for a large class of RNNs.

609, TITLE: Contamination Attacks and Mitigation in Multi-Party Machine Learning
https://papers.nips.cc/paper/7895-contamination-attacks-and-mitigation-in-multi-party-machine-learning
AUTHORS: Jamie Hayes, Olga Ohrimenko
HIGHLIGHT: We show that one needs to be careful when using this multi-party model since a potentially malicious party can taint the model by providing contaminated data.

610, TITLE: An Improved Analysis of Alternating Minimization for Structured Multi-Response Regression
https://papers.nips.cc/paper/7896-an-improved-analysis-of-alternating-minimization-for-structured-multi-response-regression
AUTHORS: Sheng Chen, Arindam Banerjee
HIGHLIGHT: In this work, we present a resampling-free analysis for the alternating minimization algorithm applied to the multi-response regression.

611, TITLE: Incorporating Context into Language Encoding Models for fMRI
https://papers.nips.cc/paper/7897-incorporating-context-into-language-encoding-models-for-fmri
AUTHORS: Shailee Jain, Alexander Huth
HIGHLIGHT: In this work we instead build encoding models using rich contextual representations derived from an LSTM language model.

612, TITLE: CatBoost: unbiased boosting with categorical features
https://papers.nips.cc/paper/7898-catboost-unbiased-boosting-with-categorical-features
AUTHORS: Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, Andrey Gulin
HIGHLIGHT: This paper presents the key algorithmic techniques behind CatBoost, a new gradient boosting toolkit.

613, TITLE: Query K-means Clustering and the Double Dixie Cup Problem
https://papers.nips.cc/paper/7899-query-k-means-clustering-and-the-double-dixie-cup-problem
AUTHORS: I Chien, Chao Pan, Olgica Milenkovic
HIGHLIGHT: We consider the problem of approximate K-means clustering with outliers and side information provided by same-cluster queries and possibly noisy answers.

614, TITLE: Training Neural Networks Using Features Replay
https://papers.nips.cc/paper/7900-training-neural-networks-using-features-replay
AUTHORS: Zhouyuan Huo, Bin Gu, Heng Huang
HIGHLIGHT: After that, we introduce features replay algorithm and prove that it is guaranteed to converge to critical points for the non-convex problem under certain conditions.

615, TITLE: Modeling Dynamic Missingness of Implicit Feedback for Recommendation
https://papers.nips.cc/paper/7901-modeling-dynamic-missingness-of-implicit-feedback-for-recommendation
AUTHORS: Menghan Wang, Mingming Gong, Xiaolin Zheng, Kun Zhang
HIGHLIGHT: To model and exploit the dynamics of missingness, we propose a latent variable named ``\emph{user intent}" to govern the temporal changes of item missingness, and a hidden Markov model to represent such a process.

616, TITLE: Representation Learning of Compositional Data
https://papers.nips.cc/paper/7902-representation-learning-of-compositional-data
AUTHORS: Marta Avalos, Richard Nock, Cheng Soon Ong, Julien Rouar, Ke Sun
HIGHLIGHT: Focusing on principal component analysis (PCA), we propose an approach that allows low dimensional representation learning directly from the original data.

617, TITLE: Model-based targeted dimensionality reduction for neuronal population data
https://papers.nips.cc/paper/7903-model-based-targeted-dimensionality-reduction-for-neuronal-population-data
AUTHORS: Mikio Aoi, Jonathan W. Pillow
HIGHLIGHT: Here we propose a new model-based method for targeted dimensionality reduction based on a probabilistic generative model of the population response data.

618, TITLE: On gradient regularizers for MMD GANs
https://papers.nips.cc/paper/7904-on-gradient-regularizers-for-mmd-gans
AUTHORS: Michael Arbel, Dougal Sutherland, Mikolaj Binkowski, Arthur Gretton
HIGHLIGHT: We propose a principled method for gradient-based regularization of the critic of GAN-like models trained by adversarially optimizing the kernel of a Maximum Mean Discrepancy (MMD).

619, TITLE: Heterogeneous Multi-output Gaussian Process Prediction
https://papers.nips.cc/paper/7905-heterogeneous-multi-output-gaussian-process-prediction
AUTHORS: Pablo Moreno-Mu?oz, Antonio Art?s, Mauricio ?lvarez
HIGHLIGHT: We present a novel extension of multi-output Gaussian processes for handling heterogeneous outputs.

620, TITLE: Large-Scale Stochastic Sampling from the Probability Simplex
https://papers.nips.cc/paper/7906-large-scale-stochastic-sampling-from-the-probability-simplex
AUTHORS: Jack Baker, Paul Fearnhead, Emily Fox, Christopher Nemeth
HIGHLIGHT: To avoid the biases caused by this discretization error, we propose the stochastic Cox-Ingersoll-Ross process (SCIR), which removes all discretization error and we prove that samples from the SCIR process are asymptotically unbiased.

621, TITLE: Policy Regret in Repeated Games
https://papers.nips.cc/paper/7907-policy-regret-in-repeated-games
AUTHORS: Raman Arora, Michael Dinitz, Teodor Vanislavov Marinov, Mehryar Mohri
HIGHLIGHT: Relating this back to external regret, we show that coarse correlated equilibria (which no-external regret players will converge to) are a strict subset of policy equilibria.

622, TITLE: A Theory-Based Evaluation of Nearest Neighbor Models Put Into Practice
https://papers.nips.cc/paper/7908-a-theory-based-evaluation-of-nearest-neighbor-models-put-into-practice
AUTHORS: Hendrik Fichtenberger, Dennis Rohde
HIGHLIGHT: In the k-nearest neighborhood model (k-NN), we are given a set of points P, and we shall answer queries q by returning the k nearest neighbors of q in P according to some metric.

623, TITLE: Banach Wasserstein GAN

https://papers.nips.cc/paper/7909-banach-wasserstein-gan
AUTHORS: Jonas Adler, Sebastian Lunz
HIGHLIGHT: We generalize the theory of WGAN with gradient penalty to Banach spaces, allowing practitioners to select the features to emphasize in the generator.

624, TITLE: Provable Gaussian Embedding with One Observation
https://papers.nips.cc/paper/7910-provable-gaussian-embedding-with-one-observation
AUTHORS: Ming Yu, Zhuoran Yang, Tuo Zhao, Mladen Kolar, Princeton Zhaoran Wang
HIGHLIGHT: In this paper, we study the Gaussian embedding model and develop the first theoretical results for exponential family embedding models.

625, TITLE: BRITS: Bidirectional Recurrent Imputation for Time Series
https://papers.nips.cc/paper/7911-brits-bidirectional-recurrent-imputation-for-time-series
AUTHORS: Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, Yitan Li
HIGHLIGHT: In this paper, we propose BRITS, a novel method based on recurrent neural networks for missing value imputation in time series data.

626, TITLE: M-Walk: Learning to Walk over Graphs using Monte Carlo Tree Search
https://papers.nips.cc/paper/7912-m-walk-learning-to-walk-over-graphs-using-monte-carlo-tree-search
AUTHORS: Yelong Shen, Jianshu Chen, Po-Sen Huang, Yuqing Guo, Jianfeng Gao
HIGHLIGHT: In order to effectively train the agent from sparse rewards, we combine MCTS with the neural policy to generate trajectories yielding more positive rewards.

627, TITLE: Extracting Relationships by Multi-Domain Matching
https://papers.nips.cc/paper/7913-extracting-relationships-by-multi-domain-matching
AUTHORS: Yitong Li, michael Murias, geraldine Dawson, David E. Carlson
HIGHLIGHT: This work proposes a novel approach, the Multiple Domain Matching Network (MDMN), to exploit this structure.
In many biological and medical contexts, we construct a large labeled corpus by aggregating many sources to use in target prediction tasks.

628, TITLE: Efficient Gradient Computation for Structured Output Learning with Rational and Tropical Losses
https://papers.nips.cc/paper/7914-efficient-gradient-computation-for-structured-output-learning-with-rational-and-tropical-losses
AUTHORS: Corinna Cortes, Vitaly Kuznetsov, Mehryar Mohri, Dmitry Storcheus, Scott Yang
HIGHLIGHT: In this paper, we design efficient gradient computation algorithms for two broad families of structured prediction loss functions: rational and tropical losses.

629, TITLE: Generative Probabilistic Novelty Detection with Adversarial Autoencoders
https://papers.nips.cc/paper/7915-generative-probabilistic-novelty-detection-with-adversarial-autoencoders
AUTHORS: Stanislav Pidhorskyi, Ranya Almohsen, Gianfranco Doretto
HIGHLIGHT: We assume that training data is available to describe only the inlier distribution.

630, TITLE: Diminishing Returns Shape Constraints for Interpretability and Regularization
https://papers.nips.cc/paper/7916-diminishing-returns-shape-constraints-for-interpretability-and-regularization
AUTHORS: Maya Gupta, Dara Bahri, Andrew Cotter, Kevin Canini
HIGHLIGHT: We investigate machine learning models that can provide diminishing returns and accelerating returns guarantees to capture prior knowledge or policies about how outputs should depend on inputs.

631, TITLE: Scalable Hyperparameter Transfer Learning
https://papers.nips.cc/paper/7917-scalable-hyperparameter-transfer-learning
AUTHORS: Valerio Perrone, Rodolphe Jenatton, Matthias W. Seeger, Cedric Archambeau
HIGHLIGHT: We propose a multi-task adaptive Bayesian linear regression model for transfer learning in BO, whose complexity is linear in the function evaluations: one Bayesian linear regression model is associated to each black-box function optimization problem (or task), while transfer learning is achieved by coupling the models through a shared deep neural net.

632, TITLE: Stochastic Nonparametric Event-Tensor Decomposition
https://papers.nips.cc/paper/7918-stochastic-nonparametric-event-tensor-decomposition
AUTHORS: Shandian Zhe, Yishuai Du
HIGHLIGHT: To address these issues, we formulate event-tensors, to preserve the complete temporal information for multiway data, and propose a novel Bayesian nonparametric decomposition model.

633, TITLE: Scaling Gaussian Process Regression with Derivatives
https://papers.nips.cc/paper/7919-scaling-gaussian-process-regression-with-derivatives
AUTHORS: David Eriksson, Kun Dong, Eric Lee, David Bindel, Andrew G. Wilson
HIGHLIGHT: We propose iterative solvers using fast $\mathcal{O}(nd)$ matrix-vector multiplications (MVMs), together with pivoted Cholesky preconditioning that cuts the iterations to convergence by several orders of magnitude, allowing for fast kernel learning and prediction.

634, TITLE: Differentially Private Testing of Identity and Closeness of Discrete Distributions
https://papers.nips.cc/paper/7920-differentially-private-testing-of-identity-and-closeness-of-discrete-distributions
AUTHORS: Jayadev Acharya, Ziteng Sun, Huanyu Zhang
HIGHLIGHT: In this work, we derive upper and lower bounds on the sample complexity of both the problems under (ε, δ)-differential privacy.

635, TITLE: Bayesian Adversarial Learning
https://papers.nips.cc/paper/7921-bayesian-adversarial-learning
AUTHORS: Nanyang Ye, Zhanxing Zhu
HIGHLIGHT: In this work, a novel robust training framework is proposed to alleviate this issue, Bayesian Robust Learning, in which a distribution is put on the adversarial data-generating distribution to account for the uncertainty of the adversarial data-generating process.

636, TITLE: Efficient Convex Completion of Coupled Tensors using Coupled Nuclear Norms
https://papers.nips.cc/paper/7922-efficient-convex-completion-of-coupled-tensors-using-coupled-nuclear-norms
AUTHORS: Kishan Wimalawarne, Hiroshi Mamitsuka
HIGHLIGHT: In this paper, we introduce a new set of coupled norms known as coupled nuclear norms by constraining the CP rank of coupled tensors.

637, TITLE: Maximizing Induced Cardinality Under a Determinantal Point Process
https://papers.nips.cc/paper/7923-maximizing-induced-cardinality-under-a-determinantal-point-process
AUTHORS: Jennifer A. Gillenwater, Alex Kulesza, Sergei Vassilvitskii, Zelda E. Mariet
HIGHLIGHT: In this paper we advocate an alternative framework for applying DPPs to recommender systems.

638, TITLE: Causal Inference with Noisy and Missing Covariates via Matrix Factorization
https://papers.nips.cc/paper/7924-causal-inference-with-noisy-and-missing-covariates-via-matrix-factorization
AUTHORS: Nathan Kallus, Xiaojie Mao, Madeleine Udell
HIGHLIGHT: We propose the use of matrix factorization to infer the confounders from noisy covariates.

639, TITLE: rho-POMDPs have Lipschitz-Continuous epsilon-Optimal Value Functions
https://papers.nips.cc/paper/7925-rho-pomdps-have-lipschitz-continuous-epsilon-optimal-value-functions
AUTHORS: Mathieu Fehr, Olivier Buffet, Vincent Thomas, Jilles Dibangoye
HIGHLIGHT: In this paper, we focus on POMDPs and ?-POMDPs with ? ? -Lipschitz reward function, and demonstrate that, for finite horizons, the optimal value function is Lipschitz-continuous.

640, TITLE: Online Structure Learning for Feed-Forward and Recurrent Sum-Product Networks
https://papers.nips.cc/paper/7926-online-structure-learning-for-feed-forward-and-recurrent-sum-product-networks
AUTHORS: Agastya Kalra, Abdullah Rashwan, Wei-Shou Hsu, Pascal Poupart, Prashant Doshi, Georgios Trimponias
HIGHLIGHT: This paper describes a new online structure learning technique for feed-forward and recurrent SPNs.

641, TITLE: Uncertainty Sampling is Preconditioned Stochastic Gradient Descent on Zero-One Loss
https://papers.nips.cc/paper/7927-uncertainty-sampling-is-preconditioned-stochastic-gradient-descent-on-zero-one-loss
AUTHORS: Stephen Mussmann, Percy S. Liang
HIGHLIGHT: In this work, we give a theoretical explanation of this phenomenon, showing that uncertainty sampling on a convex (e.g., logistic) loss can be interpreted as performing a preconditioned stochastic gradient step on the population zero-one loss.

642, TITLE: A Probabilistic U-Net for Segmentation of Ambiguous Images
https://papers.nips.cc/paper/7928-a-probabilistic-u-net-for-segmentation-of-ambiguous-images
AUTHORS: Simon Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R. Ledsam, Klaus Maier-Hein, S. M. Ali Eslami, Danilo Jimenez Rezende, Olaf Ronneberger
HIGHLIGHT: To this end we propose a generative segmentation model based on a combination of a U-Net with a conditional variational autoencoder that is capable of efficiently producing an unlimited number of plausible hypotheses.

643, TITLE: Unorganized Malicious Attacks Detection
https://papers.nips.cc/paper/7929-unorganized-malicious-attacks-detection
AUTHORS: Ming Pang, Wei Gao, Min Tao, Zhi-Hua Zhou
HIGHLIGHT: We formulate the unorganized malicious attacks detection as a matrix completion problem, and propose the Unorganized Malicious Attacks detection (UMA) algorithm, based on the alternating splitting augmented Lagrangian method.

644, TITLE: Causal Inference via Kernel Deviance Measures
https://papers.nips.cc/paper/7930-causal-inference-via-kernel-deviance-measures
AUTHORS: Jovana Mitrovic, Dino Sejdinovic, Yee Whye Teh
HIGHLIGHT: In this paper, we propose Kernel Conditional Deviance for Causal Inference (KCDC) a fully nonparametric causal discovery method based on purely observational data.

645, TITLE: Bayesian Alignments of Warped Multi-Output Gaussian Processes
https://papers.nips.cc/paper/7931-bayesian-alignments-of-warped-multi-output-gaussian-processes
AUTHORS: Markus Kaiser, Clemens Otte, Thomas Runkler, Carl Henrik Ek
HIGHLIGHT: We propose a novel Bayesian approach to modelling nonlinear alignments of time series based on latent shared information.

646, TITLE: Hybrid Macro/Micro Level Backpropagation for Training Deep Spiking Neural Networks
https://papers.nips.cc/paper/7932-hybrid-macromicro-level-backpropagation-for-training-deep-spiking-neural-networks
AUTHORS: Yingyezhe Jin, Wenrui Zhang, Peng Li
HIGHLIGHT: We present a hybrid macro/micro level backpropagation (HM2-BP) algorithm for training multi-layer SNNs.

647, TITLE: Gen-Oja: Simple & Efficient Algorithm for Streaming Generalized Eigenvector Computation
https://papers.nips.cc/paper/7933-gen-oja-simple-efficient-algorithm-for-streaming-generalized-eigenvector-computation
AUTHORS: Kush Bhatia, Aldo Pacchiano, Nicolas Flammarion, Peter L. Bartlett, Michael I. Jordan
HIGHLIGHT: In this paper, we study the problems of principle Generalized Eigenvector computation and Canonical Correlation Analysis in the stochastic setting.

648, TITLE: Efficient online algorithms for fast-rate regret bounds under sparsity
https://papers.nips.cc/paper/7934-efficient-online-algorithms-for-fast-rate-regret-bounds-under-sparsity
AUTHORS: Pierre Gaillard, Olivier Wintenberger
HIGHLIGHT: We consider the problem of online convex optimization in two different settings: arbitrary and i.i.d. sequence of convex loss functions.

649, TITLE: GILBO: One Metric to Measure Them All
https://papers.nips.cc/paper/7935-gilbo-one-metric-to-measure-them-all
AUTHORS: Alexander A. Alemi, Ian Fischer
HIGHLIGHT: We propose a simple, tractable lower bound on the mutual information contained in the joint generative density of any latent variable generative model: the GILBO (Generative Information Lower BOund).

650, TITLE: Predictive Uncertainty Estimation via Prior Networks
https://papers.nips.cc/paper/7936-predictive-uncertainty-estimation-via-prior-networks
AUTHORS: Andrey Malinin, Mark Gales
HIGHLIGHT: This work proposes a new framework for modeling predictive uncertainty called Prior Networks (PNs) which explicitly models \emph{distributional uncertainty}.

651, TITLE: Dual Policy Iteration
https://papers.nips.cc/paper/7937-dual-policy-iteration
AUTHORS: Wen Sun, Geoffrey J. Gordon, Byron Boots, J. Bagnell
HIGHLIGHT: In this work we study this Dual Policy Iteration (DPI) strategy in an alternating optimization framework and provide a convergence analysis that extends existing API theory.

652, TITLE: A probabilistic population code based on neural samples
https://papers.nips.cc/paper/7938-a-probabilistic-population-code-based-on-neural-samples
AUTHORS: Sabyasachi Shivkumar, Richard Lange, Ankani Chattoraj, Ralf Haefner
HIGHLIGHT: A central debate concerns the question of whether neural responses represent samples of latent variables (Hoyer & Hyvarinnen 2003) or parameters of their distributions (Ma et al. 2006) with efforts being made to distinguish between them (Grabska-Barwinska et al. 2013).

653, TITLE: Manifold-tiling Localized Receptive Fields are Optimal in Similarity-preserving Neural Networks
https://papers.nips.cc/paper/7939-manifold-tiling-localized-receptive-fields-are-optimal-in-similarity-preserving-neural-networks
AUTHORS: Anirvan Sengupta, Cengiz Pehlevan, Mariano Tepper, Alexander Genkin, Dmitri Chklovskii
HIGHLIGHT: Here, we propose that localized receptive fields emerge in similarity-preserving networks of rectifying neurons that learn low-dimensional manifolds populated by sensory inputs.

654, TITLE: On the Convergence and Robustness of Training GANs with Regularized Optimal Transport
https://papers.nips.cc/paper/7940-on-the-convergence-and-robustness-of-training-gans-with-regularized-optimal-transport
AUTHORS: Maziar Sanjabi, Jimmy Ba, Meisam Razaviyayn, Jason D. Lee
HIGHLIGHT: In this work, we show that obtaining gradient information of the smoothed Wasserstein GAN formulation, which is based on regularized Optimal Transport (OT), is computationally effortless and hence one can apply first order optimization methods to minimize this objective.

655, TITLE: Model-Agnostic Private Learning
https://papers.nips.cc/paper/7941-model-agnostic-private-learning
AUTHORS: Raef Bassily, Abhradeep Guha Thakurta, Om Dipakbhai Thakkar
HIGHLIGHT: We design differentially private learning algorithms that are agnostic to the learning model assuming access to limited amount of unlabeled public data.

656, TITLE: Constrained Generation of Semantically Valid Graphs via Regularizing Variational Autoencoders
https://papers.nips.cc/paper/7942-constrained-generation-of-semantically-valid-graphs-via-regularizing-variational-autoencoders
AUTHORS: Tengfei Ma, Jie Chen, Cao Xiao
HIGHLIGHT: In this work, we propose a regularization framework for variational autoencoders as a step toward semantic validity.

657, TITLE: Provably Correct Automatic Sub-Differentiation for Qualified Programs
https://papers.nips.cc/paper/7943-provably-correct-automatic-sub-differentiation-for-qualified-programs
AUTHORS: Sham M. Kakade, Jason D. Lee
HIGHLIGHT: This work considers the question: is there a \emph{Cheap Sub-gradient Principle}?

658, TITLE: Deep Homogeneous Mixture Models: Representation, Separation, and Approximation
https://papers.nips.cc/paper/7944-deep-homogeneous-mixture-models-representation-separation-and-approximation
AUTHORS: Priyank Jaini, Pascal Poupart, Yaoliang Yu
HIGHLIGHT: In this work, we formally establish the relationships among latent tree graphical models (including special cases such as hidden Markov models and tensorial mixture models), hierarchical tensor formats and sum-product networks.

659, TITLE: Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks
https://papers.nips.cc/paper/7945-parameters-as-interacting-particles-long-time-convergence-and-asymptotic-error-scaling-of-neural-networks
AUTHORS: Grant Rotskoff, Eric Vanden-Eijnden
HIGHLIGHT: In this limit, we prove a dynamical variant of the universal approximation theorem showing that the optimal representation can be attained by stochastic gradient descent, the algorithm ubiquitously used for parameter optimization in machine learning.

660, TITLE: Hierarchical Reinforcement Learning for Zero-shot Generalization with Subtask Dependencies
https://papers.nips.cc/paper/7946-hierarchical-reinforcement-learning-for-zero-shot-generalization-with-subtask-dependencies
AUTHORS: Sungryull Sohn, Junhyuk Oh, Honglak Lee
HIGHLIGHT: We introduce a new RL problem where the agent is required to generalize to a previously-unseen environment characterized by a subtask graph which describes a set of subtasks and their dependencies.

661, TITLE: A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks
https://papers.nips.cc/paper/7947-a-simple-unified-framework-for-detecting-out-of-distribution-samples-and-adversarial-attacks
AUTHORS: Kimin Lee, Kibok Lee, Honglak Lee, Jinwoo Shin
HIGHLIGHT: In this paper, we propose a simple yet effective method for detecting any abnormal samples, which is applicable to any pre-trained softmax neural classifier.

662, TITLE: End-to-End Differentiable Physics for Learning and Control
https://papers.nips.cc/paper/7948-end-to-end-differentiable-physics-for-learning-and-control
AUTHORS: Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, J. Zico Kolter

HIGHLIGHT: Specifically, in this paper we demonstrate how to perform backpropagation analytically through a physical simulator defined via a linear complementarity problem.

663, TITLE: BRUNO: A Deep Recurrent Model for Exchangeable Data
https://papers.nips.cc/paper/7949-bruno-a-deep-recurrent-model-for-exchangeable-data
AUTHORS: Iryna Korshunova, Jonas Degrave, Ferenc Huszar, Yarin Gal, Arthur Gretton, Joni Dambre
HIGHLIGHT: We present a novel model architecture which leverages deep learning tools to perform exact Bayesian inference on sets of high dimensional, complex observations.

664, TITLE: Stimulus domain transfer in recurrent models for large scale cortical population prediction on video
https://papers.nips.cc/paper/7950-stimulus-domain-transfer-in-recurrent-models-for-large-scale-cortical-population-prediction-on-video
AUTHORS: Fabian Sinz, Alexander S. Ecker, Paul Fahey, Edgar Walker, Erick Cobos, Emmanouil Froudarakis, Dimitri Yatsenko, Zachary Pitkow, Jacob Reimer, Andreas Tolias
HIGHLIGHT: We investigated these domain transfer properties in our model and find that our model trained on natural images is able to correctly predict the orientation tuning of neurons in responses to artificial noise stimuli.

665, TITLE: Mapping Images to Scene Graphs with Permutation-Invariant Structured Prediction
https://papers.nips.cc/paper/7951-mapping-images-to-scene-graphs-with-permutation-invariant-structured-prediction
AUTHORS: Roei Herzig, Moshiko Raboh, Gal Chechik, Jonathan Berant, Amir Globerson
HIGHLIGHT: Here we propose a design principle for such architectures that follows from a natural requirement of permutation invariance.

666, TITLE: Distributed Multi-Player Bandits - a Game of Thrones Approach
https://papers.nips.cc/paper/7952-distributed-multi-player-bandits-a-game-of-thrones-approach
AUTHORS: Ilai Bistritz, Amir Leshem
HIGHLIGHT: We present a distributed algorithm and prove that it achieves an expected sum of regrets of near-O\left(\log^{2}T\right).

667, TITLE: Efficient Loss-Based Decoding on Graphs for Extreme Classification
https://papers.nips.cc/paper/7953-efficient-loss-based-decoding-on-graphs-for-extreme-classification
AUTHORS: Itay Evron, Edward Moroshko, Koby Crammer
HIGHLIGHT: We build on a recent extreme classification framework with logarithmic time and space (LTLS), and on a general approach for error correcting output coding (ECOC) with loss-based decoding, and introduce a flexible and efficient approach accompanied by theoretical bounds.

668, TITLE: Chaining Mutual Information and Tightening Generalization Bounds
https://papers.nips.cc/paper/7954-chaining-mutual-information-and-tightening-generalization-bounds
AUTHORS: Amir Asadi, Emmanuel Abbe, Sergio Verdu
HIGHLIGHT: In this paper, we introduce a technique to combine chaining and mutual information methods, to obtain a generalization bound that is both algorithm-dependent and that exploits the dependencies between the hypotheses.

669, TITLE: Implicit Probabilistic Integrators for ODEs
https://papers.nips.cc/paper/7955-implicit-probabilistic-integrators-for-odes
AUTHORS: Onur Teymur, Han Cheng Lie, Tim Sullivan, Ben Calderhead
HIGHLIGHT: We introduce a family of implicit probabilistic integrators for initial value problems (IVPs), taking as a starting point the multistep Adams–Moulton method.

670, TITLE: Learning Attentional Communication for Multi-Agent Cooperation
https://papers.nips.cc/paper/7956-learning-attentional-communication-for-multi-agent-cooperation
AUTHORS: Jiechuan Jiang, Zongqing Lu
HIGHLIGHT: To tackle these difficulties, in this paper, we propose an attentional communication model that learns when communication is needed and how to integrate shared information for cooperative decision making.

671, TITLE: Training Deep Models Faster with Robust, Approximate Importance Sampling
https://papers.nips.cc/paper/7957-training-deep-models-faster-with-robust-approximate-importance-sampling
AUTHORS: Tyler B. Johnson, Carlos Guestrin
HIGHLIGHT: We propose a robust, approximate importance sampling procedure (RAIS) for stochastic gradient de- scent.

672, TITLE: Bandit Learning with Implicit Feedback

https://papers.nips.cc/paper/7958-bandit-learning-with-implicit-feedback
AUTHORS: Yi Qi, Qingyun Wu, Hongning Wang, Jie Tang, Maosong Sun
HIGHLIGHT: Since users' examination behavior is unobserved, we introduce latent variables to model it.

673, TITLE: Unsupervised Text Style Transfer using Language Models as Discriminators
https://papers.nips.cc/paper/7959-unsupervised-text-style-transfer-using-language-models-as-discriminators
AUTHORS: Zichao Yang, Zhiting Hu, Chris Dyer, Eric P. Xing, Taylor Berg-Kirkpatrick
HIGHLIGHT: In this paper, we propose a technique of using a target domain language model as the discriminator to provide richer, token-level feedback during the learning process.

674, TITLE: Relational recurrent neural networks
https://papers.nips.cc/paper/7960-relational-recurrent-neural-networks
AUTHORS: Adam Santoro, Ryan Faulkner, David Raposo, Jack Rae, Mike Chrzanowski, Theophane Weber, Daan Wierstra, Oriol Vinyals, Razvan Pascanu, Timothy Lillicrap
HIGHLIGHT: Relational recurrent neural networks

675, TITLE: Streaming Kernel PCA with <var>\tilde{O}(\sqrt{n})</var> Random Features
https://papers.nips.cc/paper/7961-streaming-kernel-pca-with-tildeosqrtn-random-features
AUTHORS: Enayat Ullah, Poorya Mianjy, Teodor Vanislavov Marinov, Raman Arora
HIGHLIGHT: We study the statistical and computational aspects of kernel principal component analysis using random Fourier features and show that under mild assumptions, $O(\sqrt{n} \log n)$ features suffices to achieve $O(1/\epsilon^2)$ sample complexity.

676, TITLE: REFUEL: Exploring Sparse Features in Deep Reinforcement Learning for Fast Disease Diagnosis
https://papers.nips.cc/paper/7962-refuel-exploring-sparse-features-in-deep-reinforcement-learning-for-fast-disease-diagnosis
AUTHORS: Yu-Shao Peng, Kai-Fu Tang, Hsuan-Tien Lin, Edward Chang
HIGHLIGHT: This paper proposes REFUEL, a reinforcement learning method with two techniques: {\em reward shaping} and {\em feature rebuilding}, to improve the performance of online symptom checking for disease diagnosis.

677, TITLE: Bayesian Model-Agnostic Meta-Learning
https://papers.nips.cc/paper/7963-bayesian-model-agnostic-meta-learning
AUTHORS: Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, Sungjin Ahn
HIGHLIGHT: In this paper, we propose a novel Bayesian model-agnostic meta-learning method.

678, TITLE: Disconnected Manifold Learning for Generative Adversarial Networks
https://papers.nips.cc/paper/7964-disconnected-manifold-learning-for-generative-adversarial-networks
AUTHORS: Mahyar Khayatkhoei, Maneesh K. Singh, Ahmed Elgammal
HIGHLIGHT: We conduct several experiments to illustrate the aforementioned shortcoming of GANs, its consequences in practice, and the effectiveness of our proposed modifications in alleviating these issues.

679, TITLE: Unsupervised Cross-Modal Alignment of Speech and Text Embedding Spaces
https://papers.nips.cc/paper/7965-unsupervised-cross-modal-alignment-of-speech-and-text-embedding-spaces
AUTHORS: Yu-An Chung, Wei-Hung Weng, Schrasing Tong, James Glass
HIGHLIGHT: Inspired by the success in unsupervised cross-lingual word embeddings, in this paper we target learning a cross-modal alignment between the embedding spaces of speech and text learned from corpora of their respective modalities in an unsupervised fashion.

680, TITLE: Learning Signed Determinantal Point Processes through the Principal Minor Assignment Problem
https://papers.nips.cc/paper/7966-learning-signed-determinantal-point-processes-through-the-principal-minor-assignment-problem
AUTHORS: Victor-Emmanuel Brunel
HIGHLIGHT: In this work, we consider a new class of DPP's, which we call signed DPP's, where we break the symmetry and allow attractive behaviors.

681, TITLE: Out-of-Distribution Detection using Multiple Semantic Label Representations
https://papers.nips.cc/paper/7967-out-of-distribution-detection-using-multiple-semantic-label-representations
AUTHORS: Gabi Shalev, Yossi Adi, Joseph Keshet
HIGHLIGHT: In this work, we consider the problem of out-of-distribution detection in neural networks.

682, TITLE: Stochastic Chebyshev Gradient Descent for Spectral Optimization
https://papers.nips.cc/paper/7968-stochastic-chebyshev-gradient-descent-for-spectral-optimization

AUTHORS: Insu Han, Haim Avron, Jinwoo Shin
HIGHLIGHT: In this paper, we develop unbiased stochastic gradients for spectral-sums, an important subclass of spectral functions.

683, TITLE: Revisiting <var>(\epsilon, \gamma, \tau)</var>-similarity learning for domain adaptation
https://papers.nips.cc/paper/7969-revisiting-epsilon-gamma-tau-similarity-learning-for-domain-adaptation
AUTHORS: Sofiane Dhouib, Ievgen Redko
HIGHLIGHT: In this paper, we propose to extend the theoretical analysis of similarity learning to the domain adaptation setting, a particular situation occurring when the similarity is learned and then deployed on samples following different probability distributions.

684, TITLE: How to tell when a clustering is (approximately) correct using convex relaxations
https://papers.nips.cc/paper/7970-how-to-tell-when-a-clustering-is-approximately-correct-using-convex-relaxations
AUTHORS: Marina Meila
HIGHLIGHT: We introduce the Sublevel Set (SS) method, a generic method to obtain sufficient guarantees of near-optimality and uniqueness (up to small perturbations) for a clustering.

685, TITLE: Constant Regret, Generalized Mixability, and Mirror Descent
https://papers.nips.cc/paper/7971-constant-regret-generalized-mixability-and-mirror-descent
AUTHORS: Zakaria Mhammedi, Robert C. Williamson
HIGHLIGHT: We consider the setting of prediction with expert advice; a learner makes predictions by aggregating those of a group of experts.

686, TITLE: A Bayesian Approach to Generative Adversarial Imitation Learning
https://papers.nips.cc/paper/7972-a-bayesian-approach-to-generative-adversarial-imitation-learning
AUTHORS: Wonseok Jeon, Seokin Seo, Kee-Eung Kim
HIGHLIGHT: A Bayesian Approach to Generative Adversarial Imitation Learning

687, TITLE: Plug-in Estimation in High-Dimensional Linear Inverse Problems: A Rigorous Analysis
https://papers.nips.cc/paper/7973-plug-in-estimation-in-high-dimensional-linear-inverse-problems-a-rigorous-analysis
AUTHORS: Alyson K. Fletcher, Parthe Pandit, Sundeep Rangan, Subrata Sarkar, Philip Schniter
HIGHLIGHT: Several recent works have considered combining linear least-squares estimation with a generic or plug-in ``denoiser" function that can be designed in a modular manner based on the prior knowledge about $\mathbf{x}$.

688, TITLE: Constrained Cross-Entropy Method for Safe Reinforcement Learning
https://papers.nips.cc/paper/7974-constrained-cross-entropy-method-for-safe-reinforcement-learning
AUTHORS: Min Wen, Ufuk Topcu
HIGHLIGHT: We propose a constrained cross-entropy-based method to solve this problem.

689, TITLE: Multi-Agent Generative Adversarial Imitation Learning
https://papers.nips.cc/paper/7975-multi-agent-generative-adversarial-imitation-learning
AUTHORS: Jiaming Song, Hongyu Ren, Dorsa Sadigh, Stefano Ermon
HIGHLIGHT: We propose a new framework for multi-agent imitation learning for general Markov games, where we build upon a generalized notion of inverse reinforcement learning.

690, TITLE: Adaptive Learning with Unknown Information Flows
https://papers.nips.cc/paper/7976-adaptive-learning-with-unknown-information-flows
AUTHORS: Yonatan Gur, Ahmadreza Momeni
HIGHLIGHT: In this paper we introduce a new, generalized MAB formulation in which additional information on each arm may appear arbitrarily throughout the decision horizon, and study the impact of such information flows on the achievable performance and the design of efficient decision-making policies.

691, TITLE: Forecasting Treatment Responses Over Time Using Recurrent Marginal Structural Networks
https://papers.nips.cc/paper/7977-forecasting-treatment-responses-over-time-using-recurrent-marginal-structural-networks
AUTHORS: Bryan Lim
HIGHLIGHT: Drawing inspiration from marginal structural models, a class of methods in epidemiology which use propensity weighting to adjust for time-dependent confounders, we introduce the Recurrent Marginal Structural Network - a sequence-to-sequence architecture for forecasting a patient's expected response to a series of planned treatments.

692, TITLE: Generative modeling for protein structures

https://papers.nips.cc/paper/7978-generative-modeling-for-protein-structures
AUTHORS: Namrata Anand, Possu Huang
HIGHLIGHT: Here, we apply Generative Adversarial Networks (GANs) to the task of generating protein structures, toward application in fast de novo protein design.

693, TITLE: Inference in Deep Gaussian Processes using Stochastic Gradient Hamiltonian Monte Carlo
https://papers.nips.cc/paper/7979-inference-in-deep-gaussian-processes-using-stochastic-gradient-hamiltonian-monte-carlo
AUTHORS: Marton Havasi, Jos? Miguel Hern?ndez-Lobato, Juan Jos? Murillo-Fuentes
HIGHLIGHT: In this work, we provide evidence for the non-Gaussian nature of the posterior and we apply the Stochastic Gradient Hamiltonian Monte Carlo method to generate samples.

694, TITLE: Knowledge Distillation by On-the-Fly Native Ensemble
https://papers.nips.cc/paper/7980-knowledge-distillation-by-on-the-fly-native-ensemble
AUTHORS: xu lan, Xiatian Zhu, Shaogang Gong
HIGHLIGHT: In this work, we present an On-the-fly Native Ensemble (ONE) learning strategy for one-stage online distillation.

695, TITLE: Non-Adversarial Mapping with VAEs
https://papers.nips.cc/paper/7981-non-adversarial-mapping-with-vaes
AUTHORS: Yedid Hoshen
HIGHLIGHT: We present an alternative approach which is able to achieve non-adversarial mapping using a novel form of Variational Auto-Encoder.

696, TITLE: Generalisation in humans and deep neural networks
https://papers.nips.cc/paper/7982-generalisation-in-humans-and-deep-neural-networks
AUTHORS: Robert Geirhos, Carlos R. M. Temme, Jonas Rauber, Heiko H. Sch?tt, Matthias Bethge, Felix A. Wichmann
HIGHLIGHT: We compare the robustness of humans and current convolutional deep neural networks (DNNs) on object recognition under twelve different types of image degradations.

697, TITLE: Towards Text Generation with Adversarially Learned Neural Outlines
https://papers.nips.cc/paper/7983-towards-text-generation-with-adversarially-learned-neural-outlines
AUTHORS: Sandeep Subramanian, Sai Rajeswar Mudumba, Alessandro Sordoni, Adam Trischler, Aaron C. Courville, Chris Pal
HIGHLIGHT: We propose a combination of both approaches with the goal of learning generative models of text.

698, TITLE: cpSGD: Communication-efficient and differentially-private distributed SGD
https://papers.nips.cc/paper/7984-cpsgd-communication-efficient-and-differentially-private-distributed-sgd
AUTHORS: Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X. Yu, Sanjiv Kumar, Brendan McMahan
HIGHLIGHT: To this end, we study algorithms that achieve both communication efficiency and differential privacy.

699, TITLE: GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration
https://papers.nips.cc/paper/7985-gpytorch-blackbox-matrix-matrix-gaussian-process-inference-with-gpu-acceleration
AUTHORS: Jacob Gardner, Geoff Pleiss, Kilian Q. Weinberger, David Bindel, Andrew G. Wilson
HIGHLIGHT: We present an efficient and general approach to GP inference based on Blackbox Matrix-Matrix multiplication (BBMM).

700, TITLE: Diffusion Maps for Textual Network Embedding
https://papers.nips.cc/paper/7986-diffusion-maps-for-textual-network-embedding
AUTHORS: Xinyuan Zhang, Yitong Li, Dinghan Shen, Lawrence Carin
HIGHLIGHT: We present diffusion maps for textual network embedding (DMTE), integrating global structural information of the graph to capture the semantic relatedness between texts, with a diffusion-convolution operation applied on the text inputs.

701, TITLE: Simple, Distributed, and Accelerated Probabilistic Programming
https://papers.nips.cc/paper/7987-simple-distributed-and-accelerated-probabilistic-programming
AUTHORS: Dustin Tran, Matthew W. Hoffman, Dave Moore, Christopher Suter, Srinivas Vasudevan, Alexey Radul
HIGHLIGHT: We describe a simple, low-level approach for embedding probabilistic programming in a deep learning ecosystem.

702, TITLE: VideoCapsuleNet: A Simplified Network for Action Detection
https://papers.nips.cc/paper/7988-videocapsulenet-a-simplified-network-for-action-detection

AUTHORS: Kevin Duarte, Yogesh Rawat, Mubarak Shah
HIGHLIGHT: In this work, we present a more elegant solution for action detection based on the recently developed capsule network.

703, TITLE: Rectangular Bounding Process
https://papers.nips.cc/paper/7989-rectangular-bounding-process
AUTHORS: Xuhui Fan, Bin Li, Scott SIsson
HIGHLIGHT: To avoid this problem we introduce a new parsimonious partition model -- the Rectangular Bounding Process (RBP) -- to efficiently partition multi-dimensional spaces, by employing a bounding strategy to enclose data points within rectangular bounding boxes.

704, TITLE: Improved Algorithms for Collaborative PAC Learning
https://papers.nips.cc/paper/7990-improved-algorithms-for-collaborative-pac-learning
AUTHORS: Huy Nguyen, Lydia Zakynthinou
HIGHLIGHT: In this work, we design new algorithms for both the realizable and the non-realizable setting, having sample complexity only $O(\ln (k))$ times the worst-case sample complexity for learning a single task.

705, TITLE: Sparse Attentive Backtracking: Temporal Credit Assignment Through Reminding
https://papers.nips.cc/paper/7991-sparse-attentive-backtracking-temporal-credit-assignment-through-reminding
AUTHORS: Nan Rosemary Ke, Anirudh Goyal ALIAS PARTH GOYAL, Olexa Bilaniuk, Jonathan Binas, Michael C. Mozer, Chris Pal, Yoshua Bengio
HIGHLIGHT: Based on this principle, we study a novel algorithm which only back-propagates through a few of these temporal skip connections, realized by a learned attention mechanism that associates current states with relevant past states.

706, TITLE: Communication Compression for Decentralized Training
https://papers.nips.cc/paper/7992-communication-compression-for-decentralized-training
AUTHORS: Hanlin Tang, Shaoduo Gan, Ce Zhang, Tong Zhang, Ji Liu
HIGHLIGHT: In this paper, We explore a natural question: {\em can the combination of both techniques lead to a system that is robust to both bandwidth and latency?}

707, TITLE: Depth-Limited Solving for Imperfect-Information Games
https://papers.nips.cc/paper/7993-depth-limited-solving-for-imperfect-information-games
AUTHORS: Noam Brown, Tuomas Sandholm, Brandon Amos
HIGHLIGHT: This paper introduces a principled way to conduct depth-limited solving in imperfect-information games by allowing the opponent to choose among a number of strategies for the remainder of the game at the depth limit.

708, TITLE: Training Deep Neural Networks with 8-bit Floating Point Numbers
https://papers.nips.cc/paper/7994-training-deep-neural-networks-with-8-bit-floating-point-numbers
AUTHORS: Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, Kailash Gopalakrishnan
HIGHLIGHT: Here we demonstrate, for the first time, the successful training of deep neural networks using 8-bit floating point numbers while fully maintaining the accuracy on a spectrum of deep learning models and datasets.

709, TITLE: Scalar Posterior Sampling with Applications
https://papers.nips.cc/paper/7995-scalar-posterior-sampling-with-applications
AUTHORS: Georgios Theocharous, Zheng Wen, Yasin Abbasi, Nikos Vlassis
HIGHLIGHT: We propose a practical non-episodic PSRL algorithm that unlike recent state-of-the-art PSRL algorithms uses a deterministic, model-independent episode switching schedule.

710, TITLE: Understanding Batch Normalization
https://papers.nips.cc/paper/7996-understanding-batch-normalization
AUTHORS: Nils Bjorck, Carla P. Gomes, Bart Selman, Kilian Q. Weinberger
HIGHLIGHT: In this paper we take a step towards a better understanding of BN, following an empirical approach.

711, TITLE: Adversarial Scene Editing: Automatic Object Removal from Weak Supervision
https://papers.nips.cc/paper/7997-adversarial-scene-editing-automatic-object-removal-from-weak-supervision
AUTHORS: Rakshith R. Shetty, Mario Fritz, Bernt Schiele
HIGHLIGHT: In this work, we take a step towards general scene-level image editing by developing an automatic interaction-free object removal model.

712, TITLE: Attacks Meet Interpretability: Attribute-steered Detection of Adversarial Samples

https://papers.nips.cc/paper/7998-attacks-meet-interpretability-attribute-steered-detection-of-adversarial-samples
AUTHORS: Guanhong Tao, Shiqing Ma, Yingqi Liu, Xiangyu Zhang
HIGHLIGHT: Therefore, we propose a novel adversarial sample detection technique for face recognition models, based on interpretability.

713, TITLE: On Neuronal Capacity
https://papers.nips.cc/paper/7999-on-neuronal-capacity
AUTHORS: Pierre Baldi, Roman Vershynin
HIGHLIGHT: We review known results, and derive new results, estimating the capacity of several neuronal models: linear and polynomial threshold gates, linear and polynomial threshold gates with constrained weights (binary weights, positive weights), and ReLU neurons.

714, TITLE: Breaking the Activation Function Bottleneck through Adaptive Parameterization
https://papers.nips.cc/paper/8000-breaking-the-activation-function-bottleneck-through-adaptive-parameterization
AUTHORS: Sebastian Flennerhag, Hujun Yin, John Keane, Mark Elliot
HIGHLIGHT: In this paper, we consider methods for making the feed-forward layer more flexible while preserving its basic structure.

715, TITLE: Learning Loop Invariants for Program Verification
https://papers.nips.cc/paper/8001-learning-loop-invariants-for-program-verification
AUTHORS: Xujie Si, Hanjun Dai, Mukund Raghothaman, Mayur Naik, Le Song
HIGHLIGHT: Inspired by how human experts construct loop invariants, we propose a reasoning framework Code2Inv that constructs the solution by multi-step decision making and querying an external program graph memory block.

716, TITLE: Cooperative Learning of Audio and Video Models from Self-Supervised Synchronization
https://papers.nips.cc/paper/8002-cooperative-learning-of-audio-and-video-models-from-self-supervised-synchronization
AUTHORS: Bruno Korbar, Du Tran, Lorenzo Torresani
HIGHLIGHT: In this work we leverage this connection to learn general and effective models for both audio and video analysis from self-supervised temporal synchronization.

717, TITLE: Towards Robust Interpretability with Self-Explaining Neural Networks
https://papers.nips.cc/paper/8003-towards-robust-interpretability-with-self-explaining-neural-networks
AUTHORS: David Alvarez Melis, Tommi Jaakkola
HIGHLIGHT: We propose three desiderata for explanations in general -- explicitness, faithfulness, and stability -- and show that existing methods do not satisfy them.

718, TITLE: Deep State Space Models for Time Series Forecasting
https://papers.nips.cc/paper/8004-deep-state-space-models-for-time-series-forecasting
AUTHORS: Syama Sundar Rangapuram, Matthias W. Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, Tim Januschowski
HIGHLIGHT: We present a novel approach to probabilistic time series forecasting that combines state space models with deep learning.

719, TITLE: Constrained Graph Variational Autoencoders for Molecule Design
https://papers.nips.cc/paper/8005-constrained-graph-variational-autoencoders-for-molecule-design
AUTHORS: Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, Alexander Gaunt
HIGHLIGHT: We propose a variational autoencoder model in which both encoder and decoder are graph-structured.

720, TITLE: Learning Libraries of Subroutines for Neurally?Guided Bayesian Program Induction
https://papers.nips.cc/paper/8006-learning-libraries-of-subroutines-for-neurallyguided-bayesian-program-induction
AUTHORS: Kevin Ellis, Lucas Morales, Mathias Sabl?-Meyer, Armando Solar-Lezama, Josh Tenenbaum
HIGHLIGHT: We contribute a program induction algorithm that learns a DSL while jointly training a neural network to efficiently search for programs in the learned DSL.

721, TITLE: Neural Architecture Optimization
https://papers.nips.cc/paper/8007-neural-architecture-optimization
AUTHORS: Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, Tie-Yan Liu
HIGHLIGHT: In this paper, we propose a simple and efficient method to automatic neural architecture design based on continuous optimization.

722, TITLE: Preference Based Adaptation for Learning Objectives
https://papers.nips.cc/paper/8008-preference-based-adaptation-for-learning-objectives
AUTHORS: Yao-Xiang Ding, Zhi-Hua Zhou
HIGHLIGHT: In this work, we discuss the task of objective adaptation, in which the learner iteratively adapts the learning objective to the underlying true objective based on the preference feedback from an oracle.

723, TITLE: Distributed <var>k</var>-Clustering for Data with Heavy Noise
https://papers.nips.cc/paper/8009-distributed-k-clustering-for-data-with-heavy-noise
AUTHORS: Shi Li, Xiangyu Guo
HIGHLIGHT: In this paper, we consider the k-center/median/means clustering with outliers problems (or the (k, z)-center/median/means problems) in the distributed setting.

724, TITLE: Beyond Log-concavity: Provable Guarantees for Sampling Multi-modal Distributions using Simulated Tempering Langevin Monte Carlo
https://papers.nips.cc/paper/8010-beyond-log-concavity-provable-guarantees-for-sampling-multi-modal-distributions-using-simulated-tempering-langevin-monte-carlo
AUTHORS: HOLDEN LEE, Andrej Risteski, Rong Ge
HIGHLIGHT: We address this problem by combining Langevin diffusion with simulated tempering.

725, TITLE: A General Method for Amortizing Variational Filtering
https://papers.nips.cc/paper/8011-a-general-method-for-amortizing-variational-filtering
AUTHORS: Joseph Marino, Milan Cvitkovic, Yisong Yue
HIGHLIGHT: We introduce the variational filtering EM algorithm, a simple, general-purpose method for performing variational inference in dynamical latent variable models using information from only past and present variables, i.e. filtering.

726, TITLE: A Reduction for Efficient LDA Topic Reconstruction
https://papers.nips.cc/paper/8012-a-reduction-for-efficient-lda-topic-reconstruction
AUTHORS: Matteo Almanza, Flavio Chierichetti, Alessandro Panconesi, Andrea Vattani
HIGHLIGHT: We present a novel approach for LDA (Latent Dirichlet Allocation) topic reconstruction.

727, TITLE: Cluster Variational Approximations for Structure Learning of Continuous-Time Bayesian Networks from Incomplete Data
https://papers.nips.cc/paper/8013-cluster-variational-approximations-for-structure-learning-of-continuous-time-bayesian-networks-from-incomplete-data
AUTHORS: Dominik Linzner, Heinz Koeppl
HIGHLIGHT: Inspired by recent advances in statistical physics, we present a new approximation scheme based on cluster-variational methods that significantly improves upon existing variational approximations.

728, TITLE: RenderNet: A deep convolutional network for differentiable rendering from 3D shapes
https://papers.nips.cc/paper/8014-rendernet-a-deep-convolutional-network-for-differentiable-rendering-from-3d-shapes
AUTHORS: Thu H. Nguyen-Phuoc, Chuan Li, Stephen Balaban, Yongliang Yang
HIGHLIGHT: We present RenderNet, a differentiable rendering convolutional network with a novel projection unit that can render 2D images from 3D shapes.

729, TITLE: Robust Hypothesis Testing Using Wasserstein Uncertainty Sets
https://papers.nips.cc/paper/8015-robust-hypothesis-testing-using-wasserstein-uncertainty-sets
AUTHORS: RUI GAO, Liyan Xie, Yao Xie, Huan Xu
HIGHLIGHT: We develop a novel computationally efficient and general framework for robust hypothesis testing.

730, TITLE: Robust Detection of Adversarial Attacks by Modeling the Intrinsic Properties of Deep Neural Networks
https://papers.nips.cc/paper/8016-robust-detection-of-adversarial-attacks-by-modeling-the-intrinsic-properties-of-deep-neural-networks
AUTHORS: Zhihao Zheng, Pengyu Hong
HIGHLIGHT: We propose an unsupervised learning approach to detect adversarial inputs without any knowledge of attackers.

731, TITLE: Monte-Carlo Tree Search for Constrained POMDPs
https://papers.nips.cc/paper/8017-monte-carlo-tree-search-for-constrained-pomdps
AUTHORS: Jongmin Lee, Geon-hyeong Kim, Pascal Poupart, Kee-Eung Kim
HIGHLIGHT: In this paper, we present CC-POMCP (Cost-Constrained POMCP), an online MCTS algorithm for large CPOMDPs that leverages the optimization of LP-induced parameters and only requires a black-box simulator of the environment.

732, TITLE: Learning to Repair Software Vulnerabilities with Generative Adversarial Networks
https://papers.nips.cc/paper/8018-learning-to-repair-software-vulnerabilities-with-generative-adversarial-networks
AUTHORS: Jacob Harer, Onur Ozdemir, Tomo Lazovich, Christopher Reale, Rebecca Russell, Louis Kim, peter chin
HIGHLIGHT: Motivated by the problem of automated repair of software vulnerabilities, we propose an adversarial learning approach that maps from one discrete source domain to another target domain without requiring paired labeled examples or source and target domains to be bijections.

733, TITLE: Layer-Wise Coordination between Encoder and Decoder for Neural Machine Translation
https://papers.nips.cc/paper/8019-layer-wise-coordination-between-encoder-and-decoder-for-neural-machine-translation
AUTHORS: Tianyu He, Xu Tan, Yingce Xia, Di He, Tao Qin, Zhibo Chen, Tie-Yan Liu
HIGHLIGHT: In this paper, we propose the concept of layer-wise coordination for NMT, which explicitly coordinates the learning of hidden representations of the encoder and decoder together layer by layer, gradually from low level to high level.

734, TITLE: Dirichlet belief networks for topic structure learning
https://papers.nips.cc/paper/8020-dirichlet-belief-networks-for-topic-structure-learning
AUTHORS: He Zhao, Lan Du, Wray Buntine, Mingyuan Zhou
HIGHLIGHT: Here we propose a new multi-layer generative process on word distributions of topics, where each layer consists of a set of topics and each topic is drawn from a mixture of the topics of the layer above.

735, TITLE: Stochastic Expectation Maximization with Variance Reduction
https://papers.nips.cc/paper/8021-stochastic-expectation-maximization-with-variance-reduction
AUTHORS: Jianfei Chen, Jun Zhu, Yee Whye Teh, Tong Zhang
HIGHLIGHT: In this paper, we propose a variance reduced stochastic EM (sEM-vr) algorithm inspired by variance reduced stochastic gradient descent algorithms.

736, TITLE: Submodular Maximization via Gradient Ascent: The Case of Deep Submodular Functions
https://papers.nips.cc/paper/8022-submodular-maximization-via-gradient-ascent-the-case-of-deep-submodular-functions
AUTHORS: Wenruo Bai, William Stafford Noble, Jeff A. Bilmes
HIGHLIGHT: We study the problem of maximizing deep submodular functions (DSFs) subject to a matroid constraint.

737, TITLE: The challenge of realistic music generation: modelling raw audio at scale
https://papers.nips.cc/paper/8023-the-challenge-of-realistic-music-generation-modelling-raw-audio-at-scale
AUTHORS: Sander Dieleman, Aaron van den Oord, Karen Simonyan
HIGHLIGHT: In this work, we explore autoregressive discrete autoencoders (ADAs) as a means to enable autoregressive models to capture long-range correlations in waveforms.

738, TITLE: Spectral Signatures in Backdoor Attacks
https://papers.nips.cc/paper/8024-spectral-signatures-in-backdoor-attacks
AUTHORS: Brandon Tran, Jerry Li, Aleksander Madry
HIGHLIGHT: In this paper, we identify a new property of all known backdoor attacks, which we call spectral signatures.

739, TITLE: Reward learning from human preferences and demonstrations in Atari
https://papers.nips.cc/paper/8025-reward-learning-from-human-preferences-and-demonstrations-in-atari
AUTHORS: Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, Dario Amodei
HIGHLIGHT: In this work, we combine two approaches to this problem: learning from expert demonstrations and learning from trajectory preferences.
Additionally, we investigate the fit of the reward model, present some reward hacking problems, and study the effects of noise in the human labels.

740, TITLE: Approximate Knowledge Compilation by Online Collapsed Importance Sampling
https://papers.nips.cc/paper/8026-approximate-knowledge-compilation-by-online-collapsed-importance-sampling
AUTHORS: Tal Friedman, Guy Van den Broeck
HIGHLIGHT: We introduce collapsed compilation, a novel approximate inference algorithm for discrete probabilistic graphical models.

741, TITLE: Neural Arithmetic Logic Units
https://papers.nips.cc/paper/8027-neural-arithmetic-logic-units
AUTHORS: Andrew Trask, Felix Hill, Scott E. Reed, Jack Rae, Chris Dyer, Phil Blunsom
HIGHLIGHT: To encourage more systematic numerical extrapolation, we propose an architecture that represents numerical quantities as linear activations which are manipulated using primitive arithmetic operators, controlled by learned gates.

742, TITLE: Pipe-SGD: A Decentralized Pipelined SGD Framework for Distributed Deep Net Training
https://papers.nips.cc/paper/8028-pipe-sgd-a-decentralized-pipelined-sgd-framework-for-distributed-deep-net-training
AUTHORS: Youjie Li, Mingchao Yu, Songze Li, Salman Avestimehr, Nam Sung Kim, Alexander Schwing
HIGHLIGHT: In this paper, we carefully analyze the AllReduce based setup, propose timing models which include network latency, bandwidth, cluster size and compute time, and demonstrate that a pipelined training with a width of two combines the best of both synchronous and asynchronous training.

743, TITLE: Improved Expressivity Through Dendritic Neural Networks
https://papers.nips.cc/paper/8029-improved-expressivity-through-dendritic-neural-networks
AUTHORS: Xundong Wu, Xiangwen Liu, Wei Li, Qing Wu
HIGHLIGHT: In this study, we model such local nonlinearity of dendritic trees with our dendritic neural network (DENN) structure and apply this structure to typical machine learning tasks.

744, TITLE: Efficient Anomaly Detection via Matrix Sketching
https://papers.nips.cc/paper/8030-efficient-anomaly-detection-via-matrix-sketching
AUTHORS: Vatsal Sharan, Parikshit Gopalan, Udi Wieder
HIGHLIGHT: Our main technical contribution is to prove matrix perturbation inequalities for operators arising in the computation of these measures.

745, TITLE: Learning to Specialize with Knowledge Distillation for Visual Question Answering
https://papers.nips.cc/paper/8031-learning-to-specialize-with-knowledge-distillation-for-visual-question-answering
AUTHORS: Jonghwan Mun, Kimin Lee, Jinwoo Shin, Bohyung Han
HIGHLIGHT: We present a principled algorithm to learn specialized models with knowledge distillation under a multiple choice learning (MCL) framework, where training examples are assigned dynamically to a subset of models for updating network parameters.

746, TITLE: A Lyapunov-based Approach to Safe Reinforcement Learning
https://papers.nips.cc/paper/8032-a-lyapunov-based-approach-to-safe-reinforcement-learning
AUTHORS: Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, Mohammad Ghavamzadeh
HIGHLIGHT: We define and present a method for constructing Lyapunov functions, which provide an effective way to guarantee the global safety of a behavior policy during training via a set of local linear constraints.

747, TITLE: Credit Assignment For Collective Multiagent RL With Global Rewards
https://papers.nips.cc/paper/8033-credit-assignment-for-collective-multiagent-rl-with-global-rewards
AUTHORS: Duc Thien Nguyen, Akshat Kumar, Hoong Chuin Lau
HIGHLIGHT: We develop collective actor-critic RL approaches for this setting, and address the problem of multiagent credit assignment, and computing low variance policy gradient estimates that result in faster convergence to high quality solutions.

748, TITLE: Statistical Optimality of Stochastic Gradient Descent on Hard Learning Problems through Multiple Passes
https://papers.nips.cc/paper/8034-statistical-optimality-of-stochastic-gradient-descent-on-hard-learning-problems-through-multiple-passes
AUTHORS: Loucas Pillaud-Vivien, Alessandro Rudi, Francis Bach
HIGHLIGHT: In order to define the notion of hardness and show that our predictive performances are optimal, we consider potentially infinite-dimensional models and notions typically associated to kernel methods, namely, the decay of eigenvalues of the covariance matrix of the features and the complexity of the optimal predictor as measured through the covariance matrix.

749, TITLE: Does mitigating ML's impact disparity require treatment disparity?
https://papers.nips.cc/paper/8035-does-mitigating-mls-impact-disparity-require-treatment-disparity
AUTHORS: Zachary Lipton, Julian McAuley, Alexandra Chouldechova
HIGHLIGHT: In this paper, we show that: (i) when sensitive and (nominally) nonsensitive features are correlated, DLPs will indirectly implement treatment disparity, undermining the policy desiderata they are designed to address; (ii) when group membership is partly revealed by other features, DLPs induce within-class discrimination; and (iii) in general, DLPs provide suboptimal trade-offs between accuracy and impact parity.

750, TITLE: Proximal Graphical Event Models
https://papers.nips.cc/paper/8036-proximal-graphical-event-models
AUTHORS: Debarun Bhattacharjya, Dharmashankar Subramanian, Tian Gao
HIGHLIGHT: We introduce proximal graphical event models (PGEM) as a representation of such datasets.

751, TITLE: Bayesian Control of Large MDPs with Unknown Dynamics in Data-Poor Environments
https://papers.nips.cc/paper/8037-bayesian-control-of-large-mdps-with-unknown-dynamics-in-data-poor-environments
AUTHORS: Mahdi Imani, Seyede Fatemeh Ghoreishi, Ulisses M. Braga-Neto
HIGHLIGHT: We propose a Bayesian decision making framework for control of Markov Decision Processes (MDPs) with unknown dynamics and large, possibly continuous, state, action, and parameter spaces in data-poor environments.

752, TITLE: Learning Overparameterized Neural Networks via Stochastic Gradient Descent on Structured Data
https://papers.nips.cc/paper/8038-learning-overparameterized-neural-networks-via-stochastic-gradient-descent-on-structured-data
AUTHORS: Yuanzhi Li, Yingyu Liang
HIGHLIGHT: Towards bridging this gap, we study the problem of learning a two-layer overparameterized ReLU neural network for multi-class classification via stochastic gradient descent (SGD) from random initialization.

753, TITLE: Hamiltonian Variational Auto-Encoder
https://papers.nips.cc/paper/8039-hamiltonian-variational-auto-encoder
AUTHORS: Anthony L. Caterini, Arnaud Doucet, Dino Sejdinovic
HIGHLIGHT: We show here how to optimally select reverse kernels in this setting and, by building upon Hamiltonian Importance Sampling (HIS) [17], we obtain a scheme that provides low-variance unbiased estimators of the ELBO and its gradients using the reparameterization trick.

754, TITLE: Modelling and unsupervised learning of symmetric deformable object categories
https://papers.nips.cc/paper/8040-modelling-and-unsupervised-learning-of-symmetric-deformable-object-categories
AUTHORS: James Thewlis, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: We propose a new approach to model and learn, without manual supervision, the symmetries of natural objects, such as faces or flowers, given only images as input.

755, TITLE: Graphical model inference: Sequential Monte Carlo meets deterministic approximations
https://papers.nips.cc/paper/8041-graphical-model-inference-sequential-monte-carlo-meets-deterministic-approximations
AUTHORS: Fredrik Lindsten, Jouni Helske, Matti Vihola
HIGHLIGHT: In this paper we present a way of bridging the gap between deterministic and stochastic inference.

756, TITLE: Statistical mechanics of low-rank tensor decomposition
https://papers.nips.cc/paper/8042-statistical-mechanics-of-low-rank-tensor-decomposition
AUTHORS: Jonathan Kadmon, Surya Ganguli
HIGHLIGHT: We derive Bayesian approximate message passing (AMP) algorithms for recovering arbitrarily shaped low-rank tensors buried within noise, and we employ dynamic mean field theory to precisely characterize their performance.

757, TITLE: Variational Bayesian Monte Carlo
https://papers.nips.cc/paper/8043-variational-bayesian-monte-carlo
AUTHORS: Luigi Acerbi
HIGHLIGHT: We introduce here a novel sample-efficient inference framework, Variational Bayesian Monte Carlo (VBMC).

758, TITLE: Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion
https://papers.nips.cc/paper/8044-sample-efficient-reinforcement-learning-with-stochastic-ensemble-value-expansion
AUTHORS: Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, Honglak Lee
HIGHLIGHT: We propose stochastic ensemble value expansion (STEVE), a novel model-based technique that addresses this issue.

759, TITLE: Efficient Online Portfolio with Logarithmic Regret
https://papers.nips.cc/paper/8045-efficient-online-portfolio-with-logarithmic-regret
AUTHORS: Haipeng Luo, Chen-Yu Wei, Kai Zheng
HIGHLIGHT: We study the decades-old problem of online portfolio management and propose the first algorithm with logarithmic regret that is not based on Cover's Universal Portfolio algorithm and admits much faster implementation.

760, TITLE: Algorithms and Theory for Multiple-Source Adaptation
https://papers.nips.cc/paper/8046-algorithms-and-theory-for-multiple-source-adaptation
AUTHORS: Judy Hoffman, Mehryar Mohri, Ningshan Zhang
HIGHLIGHT: We present a number of novel contributions to the multiple-source adaptation problem.

761, TITLE: Online Reciprocal Recommendation with Theoretical Performance Guarantees
https://papers.nips.cc/paper/8047-online-reciprocal-recommendation-with-theoretical-performance-guarantees

AUTHORS: Claudio Gentile, Nikos Parotsidis, Fabio Vitale
HIGHLIGHT: We initiate a rigorous theoretical investigation of the reciprocal recommendation task in a specific framework of sequential learning.

762, TITLE: The promises and pitfalls of Stochastic Gradient Langevin Dynamics
https://papers.nips.cc/paper/8048-the-promises-and-pitfalls-of-stochastic-gradient-langevin-dynamics
AUTHORS: Nicolas Brosse, Alain Durmus, Eric Moulines
HIGHLIGHT: We provide a detailed analysis of the Wasserstein distances between LMC, SGLD, SGLDFP and SGD and explicit expressions of the means and covariance matrices of their invariant distributions.

763, TITLE: How SGD Selects the Global Minima in Over-parameterized Learning: A Dynamical Stability Perspective
https://papers.nips.cc/paper/8049-how-sgd-selects-the-global-minima-in-over-parameterized-learning-a-dynamical-stability-perspective
AUTHORS: Lei Wu, Chao Ma, Weinan E
HIGHLIGHT: The question of which global minima are accessible by a stochastic gradient decent (SGD) algorithm with specific learning rate and batch size is studied from the perspective of dynamical stability.

764, TITLE: Differentiable MPC for End-to-end Planning and Control
https://papers.nips.cc/paper/8050-differentiable-mpc-for-end-to-end-planning-and-control
AUTHORS: Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, J. Zico Kolter
HIGHLIGHT: We present foundations for using Model Predictive Control (MPC) as a differentiable policy class for reinforcement learning.

765, TITLE: Bilevel learning of the Group Lasso structure
https://papers.nips.cc/paper/8051-bilevel-learning-of-the-group-lasso-structure
AUTHORS: Jordan Frecon, Saverio Salzo, Massimiliano Pontil
HIGHLIGHT: To circumvent this issue, we present a method to estimate the group structure by means of a continuous bilevel optimization problem where the data is split into training and validation sets.

766, TITLE: Constructing Unrestricted Adversarial Examples with Generative Models
https://papers.nips.cc/paper/8052-constructing-unrestricted-adversarial-examples-with-generative-models
AUTHORS: Yang Song, Rui Shu, Nate Kushman, Stefano Ermon
HIGHLIGHT: In this paper, we propose a new class of adversarial examples that are synthesized entirely from scratch using a conditional generative model, without being restricted to norm-bounded perturbations.

767, TITLE: Information-theoretic Limits for Community Detection in Network Models
https://papers.nips.cc/paper/8053-information-theoretic-limits-for-community-detection-in-network-models
AUTHORS: Chuyang Ke, Jean Honorio
HIGHLIGHT: We analyze the information-theoretic limits for the recovery of node labels in several network models.

768, TITLE: Learning Conditioned Graph Structures for Interpretable Visual Question Answering
https://papers.nips.cc/paper/8054-learning-conditioned-graph-structures-for-interpretable-visual-question-answering
AUTHORS: Will Norcliffe-Brown, Stathis Vafeias, Sarah Parisot
HIGHLIGHT: In this paper, we propose a novel graph-based approach for Visual Question Answering.

769, TITLE: Distributionally Robust Graphical Models
https://papers.nips.cc/paper/8055-distributionally-robust-graphical-models
AUTHORS: Rizal Fathony, Ashkan Rezaei, Mohammad Ali Bashiri, Xinhua Zhang, Brian Ziebart
HIGHLIGHT: We present adversarial graphical models (AGM), a distributionally robust approach for constructing a predictor that performs robustly for a class of data distributions defined using a graphical structure.

770, TITLE: Transfer Learning with Neural AutoML
https://papers.nips.cc/paper/8056-transfer-learning-with-neural-automl
AUTHORS: Catherine Wong, Neil Houlsby, Yifeng Lu, Andrea Gesmundo
HIGHLIGHT: To address this we propose Transfer Neural AutoML that uses knowledge from prior tasks to speed up network design.

771, TITLE: Stochastic Primal-Dual Method for Empirical Risk Minimization with O(1) Per-Iteration Complexity
https://papers.nips.cc/paper/8057-stochastic-primal-dual-method-for-empirical-risk-minimization-with-o1-per-iteration-complexity
AUTHORS: Conghui Tan, Tong Zhang, Shiqian Ma, Ji Liu

HIGHLIGHT: In this paper, we propose a new stochastic primal-dual method to solve this class of problems.

772, TITLE: On preserving non-discrimination when combining expert advice
https://papers.nips.cc/paper/8058-on-preserving-non-discrimination-when-combining-expert-advice
AUTHORS: Avrim Blum, Suriya Gunasekar, Thodoris Lykouris, Nati Srebro
HIGHLIGHT: We study the interplay between sequential decision making and avoiding discrimination against protected groups, when examples arrive online and do not follow distributional assumptions.

773, TITLE: Learning to Play With Intrinsically-Motivated, Self-Aware Agents
https://papers.nips.cc/paper/8059-learning-to-play-with-intrinsically-motivated-self-aware-agents
AUTHORS: Nick Haber, Damian Mrowca, Stephanie Wang, Li F. Fei-Fei, Daniel L. Yamins
HIGHLIGHT: Using a simple but ecologically naturalistic simulated environment in which an agent can move and interact with objects it sees, we propose a "world-model" network that learns to predict the dynamic consequences of the agent's actions.

774, TITLE: Scaling provable adversarial defenses
https://papers.nips.cc/paper/8060-scaling-provable-adversarial-defenses
AUTHORS: Eric Wong, Frank Schmidt, Jan Hendrik Metzen, J. Zico Kolter
HIGHLIGHT: In this paper, in an effort to scale these approaches to substantially larger models, we extend previous work in three main directly.

775, TITLE: Deep Network for the Integrated 3D Sensing of Multiple People in Natural Images
https://papers.nips.cc/paper/8061-deep-network-for-the-integrated-3d-sensing-of-multiple-people-in-natural-images
AUTHORS: Andrei Zanfir, Elisabeta Marinoiu, Mihai Zanfir, Alin-Ionut Popa, Cristian Sminchisescu
HIGHLIGHT: We present MubyNet -- a feed-forward, multitask, bottom up system for the integrated localization, as well as 3d pose and shape estimation, of multiple people in monocular images.

776, TITLE: Almost Optimal Algorithms for Linear Stochastic Bandits with Heavy-Tailed Payoffs
https://papers.nips.cc/paper/8062-almost-optimal-algorithms-for-linear-stochastic-bandits-with-heavy-tailed-payoffs
AUTHORS: Han Shao, Xiaotian Yu, Irwin King, Michael R. Lyu
HIGHLIGHT: In this paper, under a weaker assumption on noises, we study the problem of \underline{lin}ear stochastic {\underline b}andits with h{\underline e}avy-{\underline t}ailed payoffs (LinBET), where the distributions have finite moments of order $1+\epsilon$, for some $\epsilon\in (0,1]$.

777, TITLE: Data-dependent PAC-Bayes priors via differential privacy
https://papers.nips.cc/paper/8063-data-dependent-pac-bayes-priors-via-differential-privacy
AUTHORS: Gintare Karolina Dziugaite, Daniel M. Roy
HIGHLIGHT: The Probably Approximately Correct (PAC) Bayes framework (McAllester, 1999) can incorporate knowledge about the learning algorithm and (data) distribution through the use of distribution-dependent priors, yielding tighter generalization bounds on data-dependent posteriors.

778, TITLE: Deep Poisson gamma dynamical systems
https://papers.nips.cc/paper/8064-deep-poisson-gamma-dynamical-systems
AUTHORS: Dandan Guo, Bo Chen, Hao Zhang, Mingyuan Zhou
HIGHLIGHT: Using sophisticated but simple-to-implement data augmentation techniques, we derived closed-form Gibbs sampling update equations by first backward and upward propagating auxiliary latent counts, and then forward and downward sampling latent variables.

779, TITLE: Dimensionality Reduction has Quantifiable Imperfections: Two Geometric Bounds
https://papers.nips.cc/paper/8065-dimensionality-reduction-has-quantifiable-imperfections-two-geometric-bounds
AUTHORS: Kry Lui, Gavin Weiguang Ding, Ruitong Huang, Robert McCann
HIGHLIGHT: In this paper, we investigate Dimensionality reduction (DR) maps in an information retrieval setting from a quantitative topology point of view.

780, TITLE: Teaching Inverse Reinforcement Learners via Features and Demonstrations
https://papers.nips.cc/paper/8066-teaching-inverse-reinforcement-learners-via-features-and-demonstrations
AUTHORS: Luis Haug, Sebastian Tschiatschek, Adish Singla
HIGHLIGHT: In this paper, we study the problem of learning from demonstrations in the setting where this is not the case, i.e., where there is a mismatch between the worldviews of the learner and the expert.

781, TITLE: Wasserstein Distributionally Robust Kalman Filtering

https://papers.nips.cc/paper/8067-wasserstein-distributionally-robust-kalman-filtering
AUTHORS: Soroosh Shafieezadeh Abadeh, Viet Anh Nguyen, Daniel Kuhn, Peyman Mohajerin Mohajerin Esfahani
HIGHLIGHT: Using these ingredients, we introduce a distributionally robust Kalman filter that hedges against model risk.

782, TITLE: Generalisation of structural knowledge in the hippocampal-entorhinal system
https://papers.nips.cc/paper/8068-generalisation-of-structural-knowledge-in-the-hippocampal-entorhinal-system
AUTHORS: James Whittington, Timothy Muller, Shirely Mark, Caswell Barry, Tim Behrens
HIGHLIGHT: We propose that to generalise structural knowledge, the representations of the structure of the world, i.e. how entities in the world relate to each other, need to be separated from representations of the entities themselves.

783, TITLE: Graph Oracle Models, Lower Bounds, and Gaps for Parallel Stochastic Optimization
https://papers.nips.cc/paper/8069-graph-oracle-models-lower-bounds-and-gaps-for-parallel-stochastic-optimization
AUTHORS: Blake E. Woodworth, Jialei Wang, Adam Smith, Brendan McMahan, Nati Srebro
HIGHLIGHT: We suggest a general oracle-based framework that captures parallel stochastic optimization in different parallelization settings described by a dependency graph, and derive generic lower bounds in terms of this graph.

784, TITLE: Adversarial Regularizers in Inverse Problems
https://papers.nips.cc/paper/8070-adversarial-regularizers-in-inverse-problems
AUTHORS: Sebastian Lunz, Ozan ?ktem, Carola-Bibiane Sch?nlieb
HIGHLIGHT: We propose a new framework for applying data-driven approaches to inverse problems, using a neural network as a regularization functional.

785, TITLE: Clustering Redemption?Beyond the Impossibility of Kleinberg?s Axioms
https://papers.nips.cc/paper/8071-clustering-redemptionbeyond-the-impossibility-of-kleinbergs-axioms
AUTHORS: Vincent Cohen-Addad, Varun Kanade, Frederik Mallmann-Trenn
HIGHLIGHT: In this work, we take a different approach, based on the observation that the consistency axiom fails to be satisfied when the “correct” number of clusters changes.

786, TITLE: Co-teaching: Robust training of deep neural networks with extremely noisy labels
https://papers.nips.cc/paper/8072-co-teaching-robust-training-of-deep-neural-networks-with-extremely-noisy-labels
AUTHORS: Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, Masashi Sugiyama
HIGHLIGHT: Therefore in this paper, we propose a new deep learning paradigm called "Co-teaching" for combating with noisy labels.

787, TITLE: Variational Inverse Control with Events: A General Framework for Data-Driven Reward Definition
https://papers.nips.cc/paper/8073-variational-inverse-control-with-events-a-general-framework-for-data-driven-reward-definition
AUTHORS: Justin Fu, Avi Singh, Dibya Ghosh, Larry Yang, Sergey Levine
HIGHLIGHT: We propose inverse event-based control, which generalizes inverse reinforcement learning methods to cases where full demonstrations are not needed, such as when only samples of desired goal states are available.

788, TITLE: A convex program for bilinear inversion of sparse vectors
https://papers.nips.cc/paper/8074-a-convex-program-for-bilinear-inversion-of-sparse-vectors
AUTHORS: Alireza Aghasi, Ali Ahmed, Paul Hand, Babhru Joshi
HIGHLIGHT: We study the case where x and w are S1- and S2-sparse with respect to a random dictionary, with the sparse vectors satisfying an effective sparsity condition, and present a recovery guarantee that depends on the number of measurements as L > Omega(S1+S2)(log(K+N))^2.

789, TITLE: Adversarial Multiple Source Domain Adaptation
https://papers.nips.cc/paper/8075-adversarial-multiple-source-domain-adaptation
AUTHORS: Han Zhao, Shanghang Zhang, Guanhang Wu, Jos? M. F. Moura, Joao P. Costeira, Geoffrey J. Gordon
HIGHLIGHT: In this paper we propose new generalization bounds and algorithms under both classification and regression settings for unsupervised multiple source domain adaptation.

790, TITLE: Neural Tangent Kernel: Convergence and Generalization in Neural Networks
https://papers.nips.cc/paper/8076-neural-tangent-kernel-convergence-and-generalization-in-neural-networks
AUTHORS: Arthur Jacot, Franck Gabriel, Clement Hongler
HIGHLIGHT: We prove that the evolution of an ANN during training can also be described by a kernel: during gradient descent on the parameters of an ANN, the network function (which maps input vectors to output vectors) follows the so-called kernel gradient associated with a new object, which we call the Neural Tangent Kernel (NTK).

791, TITLE: Contextual Stochastic Block Models
https://papers.nips.cc/paper/8077-contextual-stochastic-block-models
AUTHORS: Yash Deshpande, Subhabrata Sen, Andrea Montanari, Elchanan Mossel
HIGHLIGHT: We provide the first information theoretical tight analysis for inference of latent community structure given a sparse graph along with high dimensional node covariates, correlated with the same latent communities.

792, TITLE: A Likelihood-Free Inference Framework for Population Genetic Data using Exchangeable Neural Networks
https://papers.nips.cc/paper/8078-a-likelihood-free-inference-framework-for-population-genetic-data-using-exchangeable-neural-networks
AUTHORS: Jeffrey Chan, Valerio Perrone, Jeffrey Spence, Paul Jenkins, Sara Mathieson, Yun Song
HIGHLIGHT: In this work, we develop an exchangeable neural network that performs summary statistic-free, likelihood-free inference.

793, TITLE: Sequential Attend, Infer, Repeat: Generative Modelling of Moving Objects
https://papers.nips.cc/paper/8079-sequential-attend-infer-repeat-generative-modelling-of-moving-objects
AUTHORS: Adam Kosiorek, Hyunjik Kim, Yee Whye Teh, Ingmar Posner
HIGHLIGHT: We present Sequential Attend, Infer, Repeat (SQAIR), an interpretable deep generative model for image sequences.

794, TITLE: Randomized Prior Functions for Deep Reinforcement Learning
https://papers.nips.cc/paper/8080-randomized-prior-functions-for-deep-reinforcement-learning
AUTHORS: Ian Osband, John Aslanides, Albin Cassirer
HIGHLIGHT: We highlight why this can be a crucial shortcoming and propose a simple remedy through addition of a randomized untrainable `prior' network to each ensemble member.

795, TITLE: Compact Representation of Uncertainty in Clustering
https://papers.nips.cc/paper/8081-compact-representation-of-uncertainty-in-clustering
AUTHORS: Craig Greenberg, Nicholas Monath, Ari Kobren, Patrick Flaherty, Andrew McGregor, Andrew McCallum
HIGHLIGHT: This paper presents definitions and proofs for a dynamic-programming inference procedure that computes the partition function, the marginal probability of a cluster, and the MAP clustering---all exactly.

796, TITLE: Learning without the Phase: Regularized PhaseMax Achieves Optimal Sample Complexity
https://papers.nips.cc/paper/8082-learning-without-the-phase-regularized-phasemax-achieves-optimal-sample-complexity
AUTHORS: Fariborz Salehi, Ehsan Abbasi, Babak Hassibi
HIGHLIGHT: To this end, inspired by the PhaseMax algorithm, we formulate a convex optimization problem, where the objective function relies on an initial estimate of the true signal and also includes an additive regularization term to encourage structure.

797, TITLE: Multilingual Anchoring: Interactive Topic Modeling and Alignment Across Languages
https://papers.nips.cc/paper/8083-multilingual-anchoring-interactive-topic-modeling-and-alignment-across-languages
AUTHORS: Michelle Yuan, Benjamin Van Durme, Jordan L. Ying
HIGHLIGHT: Multilingual topic models can reveal patterns in cross-lingual document collections.

798, TITLE: Estimators for Multivariate Information Measures in General Probability Spaces
https://papers.nips.cc/paper/8084-estimators-for-multivariate-information-measures-in-general-probability-spaces
AUTHORS: Arman Rahimzamani, Himanshu Asnani, Pramod Viswanath, Sreeram Kannan
HIGHLIGHT: In this paper, we define a general graph divergence measure ($\mathbb{GDM}$), generalizing the aforementioned information measures and we construct a novel estimator via a coupling trick that directly estimates these multivariate information measures using the Radon-Nikodym derivative.

799, TITLE: DeepPINK: reproducible feature selection in deep neural networks
https://papers.nips.cc/paper/8085-deeppink-reproducible-feature-selection-in-deep-neural-networks
AUTHORS: Yang Lu, Yingying Fan, Jinchi Lv, William Stafford Noble
HIGHLIGHT: In this paper, we describe a method to increase the interpretability and reproducibility of DNNs by incorporating the idea of feature selection with controlled error rate.

800, TITLE: HOUDINI: Lifelong Learning as Program Synthesis
https://papers.nips.cc/paper/8086-houdini-lifelong-learning-as-program-synthesis
AUTHORS: Lazar Valkov, Dipak Chaudhari, Akash Srivastava, Charles Sutton, Swarat Chaudhuri
HIGHLIGHT: We present a neurosymbolic framework for the lifelong learning of algorithmic tasks that mix perception and procedural reasoning.

801, TITLE: Searching for Efficient Multi-Scale Architectures for Dense Image Prediction
https://papers.nips.cc/paper/8087-searching-for-efficient-multi-scale-architectures-for-dense-image-prediction
AUTHORS: Liang-Chieh Chen, Maxwell Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, Jon Shlens
HIGHLIGHT: In this work we explore the construction of meta-learning techniques for dense image prediction focused on the tasks of scene parsing, person-part segmentation, and semantic image segmentation.

802, TITLE: Orthogonally Decoupled Variational Gaussian Processes
https://papers.nips.cc/paper/8088-orthogonally-decoupled-variational-gaussian-processes
AUTHORS: Hugh Salimbeni, Ching-An Cheng, Byron Boots, Marc Deisenroth
HIGHLIGHT: In this work, we propose an alternative decoupled parametrization.

803, TITLE: Dendritic cortical microcircuits approximate the backpropagation algorithm
https://papers.nips.cc/paper/8089-dendritic-cortical-microcircuits-approximate-the-backpropagation-algorithm
AUTHORS: Jo?o Sacramento, Rui Ponte Costa, Yoshua Bengio, Walter Senn
HIGHLIGHT: Here, we introduce a multilayer neuronal network model with simplified dendritic compartments in which error-driven synaptic plasticity adapts the network towards a global desired output.

804, TITLE: Learning Plannable Representations with Causal InfoGAN
https://papers.nips.cc/paper/8090-learning-plannable-representations-with-causal-infogan
AUTHORS: Thanard Kurutach, Aviv Tamar, Ge Yang, Stuart J. Russell, Pieter Abbeel
HIGHLIGHT: We focus on systems with high-dimensional observations, such as images, and propose an approach that naturally combines representation learning and planning.

805, TITLE: Uniform Convergence of Gradients for Non-Convex Learning and Optimization
https://papers.nips.cc/paper/8091-uniform-convergence-of-gradients-for-non-convex-learning-and-optimization
AUTHORS: Dylan J. Foster, Ayush Sekhari, Karthik Sridharan
HIGHLIGHT: We propose vector-valued Rademacher complexities as a simple, composable, and user-friendly tool to derive dimension-free uniform convergence bounds for gradients in non-convex learning problems.

806, TITLE: Automatic differentiation in ML: Where we are and where we should be going
https://papers.nips.cc/paper/8092-automatic-differentiation-in-ml-where-we-are-and-where-we-should-be-going
AUTHORS: Bart van Merrienboer, Olivier Breuleux, Arnaud Bergeron, Pascal Lamblin
HIGHLIGHT: Based on these insights, we introduce a new graph-based intermediate representation (IR) which specifically aims to efficiently support fully-general AD for array programming.

807, TITLE: A Bayesian Nonparametric View on Count-Min Sketch
https://papers.nips.cc/paper/8093-a-bayesian-nonparametric-view-on-count-min-sketch
AUTHORS: Diana Cai, Michael Mitzenmacher, Ryan P. Adams
HIGHLIGHT: We present a Bayesian view on the count-min sketch, using the same data structure, but providing a posterior distribution over the frequencies that characterizes the uncertainty arising from the hash-based approximation.

808, TITLE: Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels
https://papers.nips.cc/paper/8094-generalized-cross-entropy-loss-for-training-deep-neural-networks-with-noisy-labels
AUTHORS: Zhilu Zhang, Mert Sabuncu
HIGHLIGHT: Here, we present a theoretically grounded set of noise-robust loss functions that can be seen as a generalization of MAE and CCE.

809, TITLE: Loss Surfaces, Mode Connectivity, and Fast Ensembling of DNNs
https://papers.nips.cc/paper/8095-loss-surfaces-mode-connectivity-and-fast-ensembling-of-dnns
AUTHORS: Timur Garipov, Pavel Izmailov, Dmitrii Podoprikhin, Dmitry P. Vetrov, Andrew G. Wilson
HIGHLIGHT: We introduce a training procedure to discover these high-accuracy pathways between modes.

810, TITLE: Flexible neural representation for physics prediction
https://papers.nips.cc/paper/8096-flexible-neural-representation-for-physics-prediction
AUTHORS: Damian Mrowca, Chengxu Zhuang, Elias Wang, Nick Haber, Li F. Fei-Fei, Josh Tenenbaum, Daniel L. Yamins
HIGHLIGHT: Inspired by this ability, we propose a hierarchical particle-based object representation that covers a wide variety of types of three-dimensional objects, including both arbitrary rigid geometrical shapes and deformable materials.

811, TITLE: Legendre Decomposition for Tensors
https://papers.nips.cc/paper/8097-legendre-decomposition-for-tensors
AUTHORS: Mahito Sugiyama, Hiroyuki Nakahara, Koji Tsuda
HIGHLIGHT: We present a novel nonnegative tensor decomposition method, called Legendre decomposition, which factorizes an input tensor into a multiplicative combination of parameters.

812, TITLE: Reinforcement Learning of Theorem Proving
https://papers.nips.cc/paper/8098-reinforcement-learning-of-theorem-proving
AUTHORS: Cezary Kaliszyk, Josef Urban, Henryk Michalewski, Miroslav Ol??k
HIGHLIGHT: We introduce a theorem proving algorithm that uses practically no domain heuristics for guiding its connection-style proof search.

813, TITLE: Data Amplification: A Unified and Competitive Approach to Property Estimation
https://papers.nips.cc/paper/8099-data-amplification-a-unified-and-competitive-approach-to-property-estimation
AUTHORS: Yi HAO, Alon Orlitsky, Ananda Theertha Suresh, Yihong Wu
HIGHLIGHT: We design the first unified, linear-time, competitive, property estimator that for a wide class of properties and for all underlying distributions uses just 2n samples to achieve the performance attained by the empirical estimator with n\sqrt{\log n} samples.

814, TITLE: Group Equivariant Capsule Networks
https://papers.nips.cc/paper/8100-group-equivariant-capsule-networks
AUTHORS: Jan Eric Lenssen, Matthias Fey, Pascal Libuschewski
HIGHLIGHT: We present group equivariant capsule networks, a framework to introduce guaranteed equivariance and invariance properties to the capsule network idea.

815, TITLE: Stein Variational Gradient Descent as Moment Matching
https://papers.nips.cc/paper/8101-stein-variational-gradient-descent-as-moment-matching
AUTHORS: Qiang Liu, Dilin Wang
HIGHLIGHT: We analyze the non-asymptotic properties of SVGD, showing that there exists a set of functions, which we call the Stein matching set, whose expectations are exactly estimated by any set of particles that satisfies the fixed point equation of SVGD.

816, TITLE: Differential Privacy for Growing Databases
https://papers.nips.cc/paper/8102-differential-privacy-for-growing-databases
AUTHORS: Rachel Cummings, Sara Krehbiel, Kevin A. Lai, Uthaipon Tantipongpipat
HIGHLIGHT: Our first contribution is a novel modification of the private multiplicative weights algorithm, which provides accurate analysis of exponentially many adaptive linear queries (an expressive query class including all counting queries) for a static database.

817, TITLE: Exploration in Structured Reinforcement Learning
https://papers.nips.cc/paper/8103-exploration-in-structured-reinforcement-learning
AUTHORS: Jungseul Ok, Alexandre Proutiere, Damianos Tranos
HIGHLIGHT: We address reinforcement learning problems with finite state and action spaces where the underlying MDP has some known structure that could be potentially exploited to minimize the exploration rates of suboptimal (state, action) pairs.

818, TITLE: A Statistical Recurrent Model on the Manifold of Symmetric Positive Definite Matrices
https://papers.nips.cc/paper/8104-a-statistical-recurrent-model-on-the-manifold-of-symmetric-positive-definite-matrices
AUTHORS: Rudrasis Chakraborty, Chun-Hao Yang, Xingjian Zhen, Monami Banerjee, Derek Archer, David Vaillancourt, Vikas Singh, Baba Vemuri
HIGHLIGHT: In this work, we study the setting where the data (or measurements) are ordered, longitudinal or temporal in nature and live on a Riemannian manifold -- this setting is common in a variety of problems in statistical machine learning, vision and medical imaging.

819, TITLE: Balanced Policy Evaluation and Learning
https://papers.nips.cc/paper/8105-balanced-policy-evaluation-and-learning
AUTHORS: Nathan Kallus
HIGHLIGHT: We present a new approach to the problems of evaluating and learning personalized decision policies from observational data of past contexts, decisions, and outcomes.

820, TITLE: Distributed Multitask Reinforcement Learning with Quadratic Convergence

https://papers.nips.cc/paper/8106-distributed-multitask-reinforcement-learning-with-quadratic-convergence
AUTHORS: Rasul Tutunov, Dongho Kim, Haitham Bou Ammar
HIGHLIGHT: In this paper, we improve over state-of-the-art by deriving multitask reinforcement learning from a variational inference perspective.

821, TITLE: Improving Neural Program Synthesis with Inferred Execution Traces
https://papers.nips.cc/paper/8107-improving-neural-program-synthesis-with-inferred-execution-traces
AUTHORS: Richard Shin, Illia Polosukhin, Dawn Song
HIGHLIGHT: As in other fields of AI, deep learning-based end-to-end approaches have made great advances in program synthesis.

822, TITLE: Adaptive Path-Integral Autoencoders: Representation Learning and Planning for Dynamical Systems
https://papers.nips.cc/paper/8108-adaptive-path-integral-autoencoders-representation-learning-and-planning-for-dynamical-systems
AUTHORS: Jung-Su Ha, Young-Jin Park, Hyeok-Joo Chae, Soon-Seo Park, Han-Lim Choi
HIGHLIGHT: We present a representation learning algorithm that learns a low-dimensional latent dynamical system from high-dimensional sequential raw data, e.g., video.

823, TITLE: Policy-Conditioned Uncertainty Sets for Robust Markov Decision Processes
https://papers.nips.cc/paper/8109-policy-conditioned-uncertainty-sets-for-robust-markov-decision-processes
AUTHORS: Andrea Tirinzoni, Marek Petrik, Xiangli Chen, Brian Ziebart
HIGHLIGHT: In this work, we propose non-rectangular uncertainty sets that bound marginal moments of state-action features defined over entire trajectories through a decision process.

824, TITLE: GLoMo: Unsupervised Learning of Transferable Relational Graphs
https://papers.nips.cc/paper/8110-glomo-unsupervised-learning-of-transferable-relational-graphs
AUTHORS: Zhilin Yang, Jake Zhao, Bhuwan Dhingra, Kaiming He, William W. Cohen, Ruslan R. Salakhutdinov, Yann LeCun
HIGHLIGHT: Modern deep transfer learning approaches have mainly focused on learning generic feature vectors from one task that are transferable to other tasks, such as word embeddings in language and pretrained convolutional features in vision.

825, TITLE: Online Learning of Quantum States
https://papers.nips.cc/paper/8111-online-learning-of-quantum-states
AUTHORS: Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, Ashwin Nayak
HIGHLIGHT: We give three different ways to prove our results---using convex optimization, quantum postselection, and sequential fat-shattering dimension---which have different advantages in terms of parameters and portability.

826, TITLE: Wavelet regression and additive models for irregularly spaced data
https://papers.nips.cc/paper/8112-wavelet-regression-and-additive-models-for-irregularly-spaced-data
AUTHORS: Asad Haris, Ali Shojaie, Noah Simon
HIGHLIGHT: We present a novel approach for nonparametric regression using wavelet basis functions.

827, TITLE: Inferring Latent Velocities from Weather Radar Data using Gaussian Processes
https://papers.nips.cc/paper/8113-inferring-latent-velocities-from-weather-radar-data-using-gaussian-processes
AUTHORS: Rico Angell, Daniel R. Sheldon
HIGHLIGHT: This paper presents a Gaussian process (GP) model to reconstruct high-resolution full velocity fields across the entire US.

828, TITLE: A Structured Prediction Approach for Label Ranking
https://papers.nips.cc/paper/8114-a-structured-prediction-approach-for-label-ranking
AUTHORS: Anna Korba, Alexandre Garcia, Florence d'Alch?-Buc
HIGHLIGHT: We propose to solve a label ranking problem as a structured output regression task.

829, TITLE: Efficient High Dimensional Bayesian Optimization with Additivity and Quadrature Fourier Features
https://papers.nips.cc/paper/8115-efficient-high-dimensional-bayesian-optimization-with-additivity-and-quadrature-fourier-features
AUTHORS: Mojmir Mutny, Andreas Krause
HIGHLIGHT: To make the optimization efficient and feasible, we introduce a novel deterministic Fourier Features approximation based on numerical integration with detailed analysis for the squared exponential kernel.

830, TITLE: FastGRNN: A Fast, Accurate, Stable and Tiny Kilobyte Sized Gated Recurrent Neural Network
https://papers.nips.cc/paper/8116-fastgrnn-a-fast-accurate-stable-and-tiny-kilobyte-sized-gated-recurrent-neural-network

AUTHORS: Aditya Kusupati, Manish Singh, Kush Bhatia, Ashish Kumar, Prateek Jain, Manik Varma
HIGHLIGHT: This paper develops the FastRNN and FastGRNN algorithms to address the twin RNN limitations of inaccurate training and inefficient prediction.

831, TITLE: Reversible Recurrent Neural Networks
https://papers.nips.cc/paper/8117-reversible-recurrent-neural-networks
AUTHORS: Matthew MacKay, Paul Vicol, Jimmy Ba, Roger B. Grosse
HIGHLIGHT: We extend our technique to attention-based sequence-to-sequence models, where it maintains performance while reducing activation memory cost by a factor of 5--10 in the encoder, and a factor of 10--15 in the decoder.

832, TITLE: SING: Symbol-to-Instrument Neural Generator
https://papers.nips.cc/paper/8118-sing-symbol-to-instrument-neural-generator
AUTHORS: Alexandre Defossez, Neil Zeghidour, Nicolas Usunier, Leon Bottou, Francis Bach
HIGHLIGHT: In this work, we study the more computationally efficient alternative of generating the waveform frame-by-frame with large strides.

833, TITLE: Learning Compressed Transforms with Low Displacement Rank
https://papers.nips.cc/paper/8119-learning-compressed-transforms-with-low-displacement-rank
AUTHORS: Anna Thomas, Albert Gu, Tri Dao, Atri Rudra, Christopher R?
HIGHLIGHT: We introduce a rich class of LDR matrices with more general displacement operators, and explicitly learn over both the operators and the low-rank component.

834, TITLE: Theoretical Linear Convergence of Unfolded ISTA and Its Practical Weights and Thresholds
https://papers.nips.cc/paper/8120-theoretical-linear-convergence-of-unfolded-ista-and-its-practical-weights-and-thresholds
AUTHORS: Xiaohan Chen, Jialin Liu, Zhangyang Wang, Wotao Yin
HIGHLIGHT: In this work, we study unfolded ISTA (Iterative Shrinkage Thresholding Algorithm) for sparse signal recovery.

835, TITLE: Iterative Value-Aware Model Learning
https://papers.nips.cc/paper/8121-iterative-value-aware-model-learning
AUTHORS: Amir-massoud Farahmand
HIGHLIGHT: This paper introduces a model-based reinforcement learning (MBRL) framework that incorporates the underlying decision problem in learning the transition model of the environment.

836, TITLE: Invariant Representations without Adversarial Training
https://papers.nips.cc/paper/8122-invariant-representations-without-adversarial-training
AUTHORS: Daniel Moyer, Shuyang Gao, Rob Brekelmans, Aram Galstyan, Greg Ver Steeg
HIGHLIGHT: We show that adversarial training is unnecessary and sometimes counter-productive; we instead cast invariant representation learning as a single information-theoretic objective that can be directly optimized.

837, TITLE: Robot Learning in Homes: Improving Generalization and Reducing Dataset Bias
https://papers.nips.cc/paper/8123-robot-learning-in-homes-improving-generalization-and-reducing-dataset-bias
AUTHORS: Abhinav Gupta, Adithyavairavan Murali, Dhiraj Prakashchand Gandhi, Lerrel Pinto
HIGHLIGHT: If we aim to deploy these models in unstructured visual environments like people's homes, they will be unable to cope with the mismatch in data distribution.

838, TITLE: Learning Safe Policies with Expert Guidance
https://papers.nips.cc/paper/8124-learning-safe-policies-with-expert-guidance
AUTHORS: Jessie Huang, Fa Wu, Doina Precup, Yang Cai
HIGHLIGHT: We propose a framework for ensuring safe behavior of a reinforcement learning agent when the reward function may be difficult to specify.

839, TITLE: Bayesian multi-domain learning for cancer subtype discovery from next-generation sequencing count data
https://papers.nips.cc/paper/8125-bayesian-multi-domain-learning-for-cancer-subtype-discovery-from-next-generation-sequencing-count-data
AUTHORS: Ehsan Hajiramezanali, Siamak Zamani Dadaneh, Alireza Karbalayghareh, Mingyuan Zhou, Xiaoning Qian
HIGHLIGHT: In this paper, we develop a Bayesian Multi-Domain Learning (BMDL) model that derives domain-dependent latent representations of overdispersed count data based on hierarchical negative binomial factorization for accurate cancer subtyping even if the number of samples for a specific cancer type is small.

840, TITLE: Learning SMaLL Predictors

https://papers.nips.cc/paper/8126-learning-small-predictors
AUTHORS: Vikas Garg, Ofer Dekel, Lin Xiao
HIGHLIGHT: We introduce a new framework for learning in severely resource-constrained settings.

841, TITLE: Phase Retrieval Under a Generative Prior
https://papers.nips.cc/paper/8127-phase-retrieval-under-a-generative-prior
AUTHORS: Paul Hand, Oscar Leong, Vlad Voroninski
HIGHLIGHT: We introduce an empirical risk formulation that has favorable global geometry for gradient methods, as soon as $m = O(k)$, under the model of a multilayer fully-connected neural network with random weights.

842, TITLE: Quadrature-based features for kernel approximation
https://papers.nips.cc/paper/8128-quadrature-based-features-for-kernel-approximation
AUTHORS: Marina Munkhoeva, Yermek Kapushev, Evgeny Burnaev, Ivan Oseledets
HIGHLIGHT: Based on an efficient numerical integration technique, we propose a unifying approach that reinterprets the previous random features methods and extends to better estimates of the kernel approximation.

843, TITLE: Reducing Network Agnostophobia
https://papers.nips.cc/paper/8129-reducing-network-agnostophobia
AUTHORS: Akshay Raj Dhamija, Manuel G?nther, Terrance Boult
HIGHLIGHT: In this paper, we show that both of these approaches help, but are generally insufficient when previously unseen classes are encountered.

844, TITLE: A Stein variational Newton method
https://papers.nips.cc/paper/8130-a-stein-variational-newton-method
AUTHORS: Gianluca Detommaso, Tiangang Cui, Youssef Marzouk, Alessio Spantini, Robert Scheichl
HIGHLIGHT: In this paper, we accelerate and generalize the SVGD algorithm by including second-order information, thereby approximating a Newton-like iteration in function space.

845, TITLE: Watch Your Step: Learning Node Embeddings via Graph Attention
https://papers.nips.cc/paper/8131-watch-your-step-learning-node-embeddings-via-graph-attention
AUTHORS: Sami Abu-El-Haija, Bryan Perozzi, Rami Al-Rfou, Alexander A. Alemi
HIGHLIGHT: In particular, we propose a novel attention model on the power series of the transition matrix, which guides the random walk to optimize an upstream objective.

846, TITLE: Visual Reinforcement Learning with Imagined Goals
https://papers.nips.cc/paper/8132-visual-reinforcement-learning-with-imagined-goals
AUTHORS: Ashvin V. Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, Sergey Levine
HIGHLIGHT: In this paper, we propose an algorithm that acquires such general-purpose skills by combining unsupervised representation learning and reinforcement learning of goal-conditioned policies.

847, TITLE: Deep Predictive Coding Network with Local Recurrent Processing for Object Recognition
https://papers.nips.cc/paper/8133-deep-predictive-coding-network-with-local-recurrent-processing-for-object-recognition
AUTHORS: Kuan Han, Haiguang Wen, Yizhen Zhang, Di Fu, Eugenio Culurciello, Zhongming Liu
HIGHLIGHT: Inspired by "predictive coding" - a theory in neuroscience, we develop a bi-directional and dynamic neural network with local recurrent processing, namely predictive coding network (PCN).

848, TITLE: PAC-Bayes bounds for stable algorithms with instance-dependent priors
https://papers.nips.cc/paper/8134-pac-bayes-bounds-for-stable-algorithms-with-instance-dependent-priors
AUTHORS: Omar Rivasplata, Csaba Szepesvari, John S. Shawe-Taylor, Emilio Parrado-Hernandez, Shiliang Sun
HIGHLIGHT: In this paper the PAC-Bayes approach is combined with stability of the hypothesis learned by a Hilbert space valued algorithm.

849, TITLE: Beyond Grids: Learning Graph Representations for Visual Recognition
https://papers.nips.cc/paper/8135-beyond-grids-learning-graph-representations-for-visual-recognition
AUTHORS: Yin Li, Abhinav Gupta
HIGHLIGHT: We propose learning graph representations from 2D feature maps for visual recognition.

850, TITLE: The Limit Points of (Optimistic) Gradient Descent in Min-Max Optimization
https://papers.nips.cc/paper/8136-the-limit-points-of-optimistic-gradient-descent-in-min-max-optimization
AUTHORS: Constantinos Daskalakis, Ioannis Panageas

HIGHLIGHT: We characterize the limit points of two basic first order methods, namely Gradient Descent/Ascent (GDA) and Optimistic Gradient Descent Ascent (OGDA).

851, TITLE: Coordinate Descent with Bandit Sampling
https://papers.nips.cc/paper/8137-coordinate-descent-with-bandit-sampling
AUTHORS: Farnood Salehi, Patrick Thiran, Elisa Celis
HIGHLIGHT: Therefore, we propose a new adaptive method for coordinate descent.

852, TITLE: Deep Dynamical Modeling and Control of Unsteady Fluid Flows
https://papers.nips.cc/paper/8138-deep-dynamical-modeling-and-control-of-unsteady-fluid-flows
AUTHORS: Jeremy Morton, Antony Jameson, Mykel J. Kochenderfer, Freddie Witherden
HIGHLIGHT: We present a method for learning the forced and unforced dynamics of airflow over a cylinder directly from CFD data.

853, TITLE: Confounding-Robust Policy Improvement
https://papers.nips.cc/paper/8139-confounding-robust-policy-improvement
AUTHORS: Nathan Kallus, Angela Zhou
HIGHLIGHT: We study the problem of learning personalized decision policies from observational data while accounting for possible unobserved confounding in the data-generating process.

854, TITLE: The Importance of Sampling inMeta-Reinforcement Learning
https://papers.nips.cc/paper/8140-the-importance-of-sampling-inmeta-reinforcement-learning
AUTHORS: Bradly Stadie, Ge Yang, Rein Houthooft, Peter Chen, Yan Duan, Yuhuai Wu, Pieter Abbeel, Ilya Sutskever
HIGHLIGHT: We interpret meta-reinforcement learning as the problem of learning how to quickly find a good sampling distribution in a new environment.

855, TITLE: Representer Point Selection for Explaining Deep Neural Networks
https://papers.nips.cc/paper/8141-representer-point-selection-for-explaining-deep-neural-networks
AUTHORS: Chih-Kuan Yeh, Joon Kim, Ian En-Hsu Yen, Pradeep K. Ravikumar
HIGHLIGHT: We propose to explain the predictions of a deep neural network, by pointing to the set of what we call representer points in the training set, for a given test point prediction.

856, TITLE: The Effect of Network Width on the Performance of Large-batch Training
https://papers.nips.cc/paper/8142-the-effect-of-network-width-on-the-performance-of-large-batch-training
AUTHORS: Lingjiao Chen, Hongyi Wang, Jinman Zhao, Dimitris Papailiopoulos, Paraschos Koutris
HIGHLIGHT: In this work, we take a first step towards analyzing how the structure (width and depth) of a neural network affects the performance of large-batch training.

857, TITLE: SNIPER: Efficient Multi-Scale Training
https://papers.nips.cc/paper/8143-sniper-efficient-multi-scale-training
AUTHORS: Bharat Singh, Mahyar Najibi, Larry S. Davis
HIGHLIGHT: We present SNIPER, an algorithm for performing efficient multi-scale training in instance level visual recognition tasks.

858, TITLE: The Sample Complexity of Semi-Supervised Learning with Nonparametric Mixture Models
https://papers.nips.cc/paper/8144-the-sample-complexity-of-semi-supervised-learning-with-nonparametric-mixture-models
AUTHORS: Chen Dan, Liu Leqi, Bryon Aragam, Pradeep K. Ravikumar, Eric P. Xing
HIGHLIGHT: We study the sample complexity of semi-supervised learning (SSL) and introduce new assumptions based on the mismatch between a mixture model learned from unlabeled data and the true mixture model induced by the (unknown) class conditional distributions.

859, TITLE: Hardware Conditioned Policies for Multi-Robot Transfer Learning
https://papers.nips.cc/paper/8145-hardware-conditioned-policies-for-multi-robot-transfer-learning
AUTHORS: Tao Chen, Adithyavairavan Murali, Abhinav Gupta
HIGHLIGHT: We propose a novel approach called Hardware Conditioned Policies where we train a universal policy conditioned on a vector representation of robot hardware.

860, TITLE: Co-regularized Alignment for Unsupervised Domain Adaptation
https://papers.nips.cc/paper/8146-co-regularized-alignment-for-unsupervised-domain-adaptation

AUTHORS: Abhishek Kumar, Prasanna Sattigeri, Kahini Wadhawan, Leonid Karlinsky, Rogerio Feris, Bill Freeman, Gregory Wornell
HIGHLIGHT: We propose co-regularized domain alignment for unsupervised domain adaptation, which constructs multiple diverse feature spaces and aligns source and target distributions in each of them individually, while encouraging that alignments agree with each other with regard to the class predictions on the unlabeled target examples.

861, TITLE: Statistical and Computational Trade-Offs in Kernel K-Means
https://papers.nips.cc/paper/8147-statistical-and-computational-trade-offs-in-kernel-k-means
AUTHORS: Daniele Calandriello, Lorenzo Rosasco
HIGHLIGHT: More precisely, we study a Nystr\"om approach to kernel k-means.

862, TITLE: Assessing the Scalability of Biologically-Motivated Deep Learning Algorithms and Architectures
https://papers.nips.cc/paper/8148-assessing-the-scalability-of-biologically-motivated-deep-learning-algorithms-and-architectures
AUTHORS: Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E. Hinton, Timothy Lillicrap
HIGHLIGHT: Here we present results on scaling up biologically motivated models of deep learning on datasets which need deep networks with appropriate architectures to achieve good performance.

863, TITLE: Learning Attractor Dynamics for Generative Memory
https://papers.nips.cc/paper/8149-learning-attractor-dynamics-for-generative-memory
AUTHORS: Yan Wu, Gregory Wayne, Karol Gregor, Timothy Lillicrap
HIGHLIGHT: In this work, we exploit recent advances in variational inference and avoid the vanishing gradient problem by training a generative distributed memory with a variational lower-bound-based Lyapunov function.

864, TITLE: The emergence of multiple retinal cell types through efficient coding of natural movies
https://papers.nips.cc/paper/8150-the-emergence-of-multiple-retinal-cell-types-through-efficient-coding-of-natural-movies
AUTHORS: Samuel Ocko, Jack Lindsey, Surya Ganguli, Stephane Deny
HIGHLIGHT: We find that optimizing the receptive fields and cell densities of two cell types makes them match the properties of the two main cell types in the primate retina, midget and parasol cells, in terms of spatial and temporal sensitivity, cell spacing, and their relative ratio.

865, TITLE: Gather-Excite: Exploiting Feature Context in Convolutional Neural Networks
https://papers.nips.cc/paper/8151-gather-excite-exploiting-feature-context-in-convolutional-neural-networks
AUTHORS: Jie Hu, Li Shen, Samuel Albanie, Gang Sun, Andrea Vedaldi
HIGHLIGHT: In this work, we propose a simple, lightweight approach for better context exploitation in CNNs.

866, TITLE: The Global Anchor Method for Quantifying Linguistic Shifts and Domain Adaptation
https://papers.nips.cc/paper/8152-the-global-anchor-method-for-quantifying-linguistic-shifts-and-domain-adaptation
AUTHORS: Zi Yin, Vin Sachidananda, Balaji Prabhakar
HIGHLIGHT: In this paper, we introduce the global anchor method for detecting corpus-level language shifts.

867, TITLE: Identification and Estimation of Causal Effects from Dependent Data
https://papers.nips.cc/paper/8153-identification-and-estimation-of-causal-effects-from-dependent-data
AUTHORS: Eli Sherman, Ilya Shpitser
HIGHLIGHT: In this paper we develop a general theory describing when causal inferences are possible in such scenarios.

868, TITLE: Deepcode: Feedback Codes via Deep Learning
https://papers.nips.cc/paper/8154-deepcode-feedback-codes-via-deep-learning
AUTHORS: Hyeji Kim, Yihan Jiang, Sreeram Kannan, Sewoong Oh, Pramod Viswanath
HIGHLIGHT: In this work, we present the first family of codes obtained via deep learning, which significantly beats state-of-the-art codes designed over several decades of research.

869, TITLE: Learning and Testing Causal Models with Interventions
https://papers.nips.cc/paper/8155-learning-and-testing-causal-models-with-interventions
AUTHORS: Jayadev Acharya, Arnab Bhattacharyya, Constantinos Daskalakis, Saravanan Kandasamy
HIGHLIGHT: We consider testing and learning problems on causal Bayesian networks as defined by Pearl (Pearl, 2009).

870, TITLE: Implicit Bias of Gradient Descent on Linear Convolutional Networks
https://papers.nips.cc/paper/8156-implicit-bias-of-gradient-descent-on-linear-convolutional-networks
AUTHORS: Suriya Gunasekar, Jason D. Lee, Daniel Soudry, Nati Srebro

HIGHLIGHT: We show that gradient descent on full-width linear convolutional networks of depth L converges to a linear predictor related to the $\ell_{2/L}$ bridge penalty in the frequency domain.

871, TITLE: DAGs with NO TEARS: Continuous Optimization for Structure Learning
https://papers.nips.cc/paper/8157-dags-with-no-tears-continuous-optimization-for-structure-learning
AUTHORS: Xun Zheng, Bryon Aragam, Pradeep K. Ravikumar, Eric P. Xing
HIGHLIGHT: In this paper, we introduce a fundamentally different strategy: we formulate the structure learning problem as a purely continuous optimization problem over real matrices that avoids this combinatorial constraint entirely.

872, TITLE: PAC-Bayes Tree: weighted subtrees with guarantees
https://papers.nips.cc/paper/8158-pac-bayes-tree-weighted-subtrees-with-guarantees
AUTHORS: Tin D. Nguyen, Samory Kpotufe
HIGHLIGHT: We present a weighted-majority classification approach over subtrees of a fixed tree, which provably achieves excess-risk of the same order as the best tree-pruning.

873, TITLE: Multi-objective Maximization of Monotone Submodular Functions with Cardinality Constraint
https://papers.nips.cc/paper/8159-multi-objective-maximization-of-monotone-submodular-functions-with-cardinality-constraint
AUTHORS: Rajan Udwani
HIGHLIGHT: We consider the problem of multi-objective maximization of monotone submodular functions subject to cardinality constraint, often formulated as $\max_{|A|=k}\min_{i\in\{1,\dots,m\}}f_i(A)$.

874, TITLE: Sanity Checks for Saliency Maps
https://papers.nips.cc/paper/8160-sanity-checks-for-saliency-maps
AUTHORS: Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, Been Kim
HIGHLIGHT: In this work, we propose an actionable methodology to evaluate what kinds of explanations a given method can and cannot provide.

875, TITLE: Probabilistic Model-Agnostic Meta-Learning
https://papers.nips.cc/paper/8161-probabilistic-model-agnostic-meta-learning
AUTHORS: Chelsea Finn, Kelvin Xu, Sergey Levine
HIGHLIGHT: In this paper, we propose a probabilistic meta-learning algorithm that can sample models for a new task from a model distribution.

876, TITLE: Reinforcement Learning with Multiple Experts: A Bayesian Model Combination Approach
https://papers.nips.cc/paper/8162-reinforcement-learning-with-multiple-experts-a-bayesian-model-combination-approach
AUTHORS: Michael Gimelfarb, Scott Sanner, Chi-Guhn Lee
HIGHLIGHT: In this paper, we apply Bayesian Model Combination with multiple experts in a way that learns to trust a good combination of experts as training progresses.

877, TITLE: e-SNLI: Natural Language Inference with Natural Language Explanations
https://papers.nips.cc/paper/8163-e-snli-natural-language-inference-with-natural-language-explanations
AUTHORS: Oana-Maria Camburu, Tim Rockt?schel, Thomas Lukasiewicz, Phil Blunsom
HIGHLIGHT: In this work, we extend the Stanford Natural Language Inference dataset with an additional layer of human-annotated natural language explanations of the entailment relations.

878, TITLE: Fast Approximate Natural Gradient Descent in a Kronecker Factored Eigenbasis
https://papers.nips.cc/paper/8164-fast-approximate-natural-gradient-descent-in-a-kronecker-factored-eigenbasis
AUTHORS: Thomas George, C?sar Laurent, Xavier Bouthillier, Nicolas Ballas, Pascal Vincent
HIGHLIGHT: In the present work we draw inspiration from both to propose a novel approximation that is provably better than KFAC and amendable to cheap partial updates.

879, TITLE: Learning convex bounds for linear quadratic control policy synthesis
https://papers.nips.cc/paper/8165-learning-convex-bounds-for-linear-quadratic-control-policy-synthesis
AUTHORS: Jack Umenberger, Thomas B. Sch?n
HIGHLIGHT: We present a method to optimize the expected value of the reward over the posterior distribution of the unknown system parameters, given data.

880, TITLE: Neural Proximal Gradient Descent for Compressive Imaging
https://papers.nips.cc/paper/8166-neural-proximal-gradient-descent-for-compressive-imaging

AUTHORS: Morteza Mardani, Qingyun Sun, David Donoho, Vardan Papyan, Hatef Monajemi, Shreyas Vasanawala, John Pauly
HIGHLIGHT: We develop a successful system solving all these challenges, using as basic architecture the repetitive application of alternating proximal and data fidelity constraints.

881, TITLE: Towards Understanding Learning Representations: To What Extent Do Different Neural Networks Learn the Same Representation
https://papers.nips.cc/paper/8167-towards-understanding-learning-representations-to-what-extent-do-different-neural-networks-learn-the-same-representation
AUTHORS: Liwei Wang, Lunjia Hu, Jiayuan Gu, Zhiqiang Hu, Yue Wu, Kun He, John Hopcroft
HIGHLIGHT: In this work, we move a tiny step towards a theory and better understanding of the representations.

882, TITLE: Optimal Algorithms for Continuous Non-monotone Submodular and DR-Submodular Maximization
https://papers.nips.cc/paper/8168-optimal-algorithms-for-continuous-non-monotone-submodular-and-dr-submodular-maximization
AUTHORS: Rad Niazadeh, Tim Roughgarden, Joshua Wang
HIGHLIGHT: In this paper we study the fundamental problems of maximizing a continuous non monotone submodular function over a hypercube, with and without coordinate-wise concavity.

883, TITLE: An intriguing failing of convolutional neural networks and the CoordConv solution
https://papers.nips.cc/paper/8169-an-intriguing-failing-of-convolutional-neural-networks-and-the-coordconv-solution
AUTHORS: Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, Jason Yosinski
HIGHLIGHT: In this paper we show a striking counterexample to this intuition via the seemingly trivial coordinate transform problem, which simply requires learning a mapping between coordinates in (x,y) Cartesian space and coordinates in one-hot pixel space.

884, TITLE: Trading robust representations for sample complexity through self-supervised visual experience
https://papers.nips.cc/paper/8170-trading-robust-representations-for-sample-complexity-through-self-supervised-visual-experience
AUTHORS: Andrea Tacchetti, Stephen Voinea, Georgios Evangelopoulos
HIGHLIGHT: We introduce a novel loss function for representation learning using unlabeled image sets and video sequences, and experimentally demonstrate that these representations support one-shot learning and reduce the sample complexity of multiple recognition tasks.

885, TITLE: Invertibility of Convolutional Generative Networks from Partial Measurements
https://papers.nips.cc/paper/8171-invertibility-of-convolutional-generative-networks-from-partial-measurements
AUTHORS: Fangchang Ma, Ulas Ayaz, Sertac Karaman
HIGHLIGHT: In this work, we present new theoretical results on convolutional generative neural networks, in particular their invertibility (i.e., the recovery of input latent code given the network output).

886, TITLE: Ex ante coordination and collusion in zero-sum multi-player extensive-form games
https://papers.nips.cc/paper/8172-ex-ante-coordination-and-collusion-in-zero-sum-multi-player-extensive-form-games
AUTHORS: Gabriele Farina, Andrea Celli, Nicola Gatti, Tuomas Sandholm
HIGHLIGHT: We propose a new game representation, the realization form, that generalizes the sequence form but can also be applied to imperfect-recall games.

887, TITLE: Multi-Agent Reinforcement Learning via Double Averaging Primal-Dual Optimization
https://papers.nips.cc/paper/8173-multi-agent-reinforcement-learning-via-double-averaging-primal-dual-optimization
AUTHORS: Hoi-To Wai, Zhuoran Yang, Princeton Zhaoran Wang, Mingyi Hong
HIGHLIGHT: In this paper, we propose a double averaging scheme, where each agent iteratively performs averaging over both space and time to incorporate neighboring gradient information and local reward information, respectively.

888, TITLE: Improving Online Algorithms via ML Predictions
https://papers.nips.cc/paper/8174-improving-online-algorithms-via-ml-predictions
AUTHORS: Manish Purohit, Zoya Svitkina, Ravi Kumar
HIGHLIGHT: In this work we study the problem of using machine-learned predictions to improve performance of online algorithms.

889, TITLE: Global Non-convex Optimization with Discretized Diffusions
https://papers.nips.cc/paper/8175-global-non-convex-optimization-with-discretized-diffusions
AUTHORS: Murat A. Erdogdu, Lester Mackey, Ohad Shamir
HIGHLIGHT: We show that this property holds for any suitably smooth diffusion and that different diffusions are suitable for optimizing different classes of convex and non-convex functions.

890, TITLE: Theoretical guarantees for EM under misspecified Gaussian mixture models
https://papers.nips.cc/paper/8176-theoretical-guarantees-for-em-under-misspecified-gaussian-mixture-models
AUTHORS: Raaz Dwivedi, nh?t H?, Koulik Khamaru, Martin J. Wainwright, Michael I. Jordan
HIGHLIGHT: We provide non-asymptotic guarantees for population and sample-based EM for parameter estimation under a few specific univariate settings of misspecified Gaussian mixture models.

891, TITLE: Coupled Variational Bayes via Optimization Embedding
https://papers.nips.cc/paper/8177-coupled-variational-bayes-via-optimization-embedding
AUTHORS: Bo Dai, Hanjun Dai, Niao He, Weiyang Liu, Zhen Liu, Jianshu Chen, Lin Xiao, Le Song
HIGHLIGHT: In this paper, we proposed coupled variational Bayes which exploits the primal-dual view of the ELBO with the variational distribution class generated by an optimization procedure, which is termed optimization embedding.

892, TITLE: Improving Explorability in Variational Inference with Annealed Variational Objectives
https://papers.nips.cc/paper/8178-improving-explorability-in-variational-inference-with-annealed-variational-objectives
AUTHORS: Chin-Wei Huang, Shawn Tan, Alexandre Lacoste, Aaron C. Courville
HIGHLIGHT: We demonstrate the drawbacks of biasing the true posterior to be unimodal, and introduce Annealed Variational Objectives (AVO) into the training of hierarchical variational methods.

893, TITLE: Latent Alignment and Variational Attention
https://papers.nips.cc/paper/8179-latent-alignment-and-variational-attention
AUTHORS: Yuntian Deng, Yoon Kim, Justin Chiu, Demi Guo, Alexander Rush
HIGHLIGHT: This work considers variational attention networks, alternatives to soft and hard attention for learning latent variable alignment models, with tighter approximation bounds based on amortized variational inference.

894, TITLE: Towards Deep Conversational Recommendations
https://papers.nips.cc/paper/8180-towards-deep-conversational-recommendations
AUTHORS: Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, Chris Pal
HIGHLIGHT: In particular we explore new neural architectures, mechanisms and methods suitable for composing conversational recommendation systems.

895, TITLE: Unsupervised Depth Estimation, 3D Face Rotation and Replacement
https://papers.nips.cc/paper/8181-unsupervised-depth-estimation-3d-face-rotation-and-replacement
AUTHORS: Joel Ruben Antony Moniz, Christopher Beckham, Simon Rajotte, Sina Honari, Chris Pal
HIGHLIGHT: We present an unsupervised approach for learning to estimate three dimensional (3D) facial structure from a single image while also predicting 3D viewpoint transformations that match a desired pose and facial geometry.

896, TITLE: Generalization Bounds for Uniformly Stable Algorithms
https://papers.nips.cc/paper/8182-generalization-bounds-for-uniformly-stable-algorithms
AUTHORS: Vitaly Feldman, Jan Vondrak
HIGHLIGHT: Here we prove substantially stronger generalization bounds for uniformly stable algorithms without any additional assumptions.

897, TITLE: Deep Anomaly Detection Using Geometric Transformations
https://papers.nips.cc/paper/8183-deep-anomaly-detection-using-geometric-transformations
AUTHORS: Izhak Golan, Ran El-Yaniv
HIGHLIGHT: We consider the problem of anomaly detection in images, and present a new detection technique.

898, TITLE: Large Scale computation of Means and Clusters for Persistence Diagrams using Optimal Transport
https://papers.nips.cc/paper/8184-large-scale-computation-of-means-and-clusters-for-persistence-diagrams-using-optimal-transport
AUTHORS: Theo Lacombe, Marco Cuturi, Steve OUDOT
HIGHLIGHT: We propose in this article a tractable framework to carry out standard tasks on PDs at scale, notably evaluating distances, estimating barycenters and performing clustering.

899, TITLE: Entropy Rate Estimation for Markov Chains with Large State Space
https://papers.nips.cc/paper/8185-entropy-rate-estimation-for-markov-chains-with-large-state-space
AUTHORS: Yanjun Han, Jiantao Jiao, Chuan-Zheng Lee, Tsachy Weissman, Yihong Wu, Tiancheng Yu
HIGHLIGHT: Extending the theory and algorithms for entropy estimation to dependent data, this paper considers the problem of estimating the entropy rate of a stationary reversible Markov chain with S states from a sample path of n observations.

900, TITLE: Adaptive Methods for Nonconvex Optimization
https://papers.nips.cc/paper/8186-adaptive-methods-for-nonconvex-optimization
AUTHORS: Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, Sanjiv Kumar
HIGHLIGHT: In this work, we provide a new analysis of such methods applied to nonconvex stochastic optimization problems, characterizing the effect of increasing minibatch size.

901, TITLE: Object-Oriented Dynamics Predictor
https://papers.nips.cc/paper/8187-object-oriented-dynamics-predictor
AUTHORS: Guangxiang Zhu, Zhiao Huang, Chongjie Zhang
HIGHLIGHT: In this paper, we present a novel object-oriented framework, called object-oriented dynamics predictor (OODP), which decomposes the environment into objects and predicts the dynamics of objects conditioned on both actions and object-to-object relations.

902, TITLE: Adaptive Skip Intervals: Temporal Abstraction for Recurrent Dynamical Models
https://papers.nips.cc/paper/8188-adaptive-skip-intervals-temporal-abstraction-for-recurrent-dynamical-models
AUTHORS: Alexander Neitz, Giambattista Parascandolo, Stefan Bauer, Bernhard Sch?lkopf
HIGHLIGHT: We introduce a method which enables a recurrent dynamics model to be temporally abstract.

903, TITLE: Scalable End-to-End Autonomous Vehicle Testing via Rare-event Simulation
https://papers.nips.cc/paper/8189-scalable-end-to-end-autonomous-vehicle-testing-via-rare-event-simulation
AUTHORS: Matthew O'Kelly, Aman Sinha, Hongseok Namkoong, Russ Tedrake, John C. Duchi
HIGHLIGHT: We implement a simulation framework that can test an entire modern autonomous driving system, including, in particular, systems that employ deep-learning perception and control algorithms.

904, TITLE: Reinforcement Learning for Solving the Vehicle Routing Problem
https://papers.nips.cc/paper/8190-reinforcement-learning-for-solving-the-vehicle-routing-problem
AUTHORS: MohammadReza Nazari, Afshin Oroojlooy, Lawrence Snyder, Martin Takac
HIGHLIGHT: We present an end-to-end framework for solving the Vehicle Routing Problem (VRP) using reinforcement learning.

905, TITLE: ATOMO: Communication-efficient Learning via Atomic Sparsification
https://papers.nips.cc/paper/8191-atomo-communication-efficient-learning-via-atomic-sparsification
AUTHORS: Hongyi Wang, Scott Sievert, Shengchao Liu, Zachary Charles, Dimitris Papailiopoulos, Stephen Wright
HIGHLIGHT: We present ATOMO, a general framework for atomic sparsification of stochastic gradients.

906, TITLE: Dynamic Network Model from Partial Observations
https://papers.nips.cc/paper/8192-dynamic-network-model-from-partial-observations
AUTHORS: Elahe Ghalebi, Baharan Mirzasoleiman, Radu Grosu, Jure Leskovec
HIGHLIGHT: We propose a novel framework for providing a non-parametric dynamic network model---based on a mixture of coupled hierarchical Dirichlet processes---based on data capturing cascade node infection times.

907, TITLE: Life-Long Disentangled Representation Learning with Cross-Domain Latent Homologies
https://papers.nips.cc/paper/8193-life-long-disentangled-representation-learning-with-cross-domain-latent-homologies
AUTHORS: Alessandro Achille, Tom Eccles, Loic Matthey, Chris Burgess, Nicholas Watters, Alexander Lerchner, Irina Higgins
HIGHLIGHT: We propose a novel algorithm for unsupervised representation learning from piece-wise stationary visual data: Variational Autoencoder with Shared Embeddings (VASE).

908, TITLE: Maximizing acquisition functions for Bayesian optimization
https://papers.nips.cc/paper/8194-maximizing-acquisition-functions-for-bayesian-optimization
AUTHORS: James Wilson, Frank Hutter, Marc Deisenroth
HIGHLIGHT: Subsequently, we identify a common family of acquisition functions, including EI and UCB, whose characteristics not only facilitate but justify use of greedy approaches for their maximization.

909, TITLE: On Markov Chain Gradient Descent
https://papers.nips.cc/paper/8195-on-markov-chain-gradient-descent
AUTHORS: Tao Sun, Yuejiao Sun, Wotao Yin
HIGHLIGHT: To obtain these results, we introduce a new technique that varies the mixing levels of the Markov chains.

910, TITLE: Variance-Reduced Stochastic Gradient Descent on Streaming Data
https://papers.nips.cc/paper/8196-variance-reduced-stochastic-gradient-descent-on-streaming-data
AUTHORS: Ellango Jothimurugesan, Ashraf Tahmasbi, Phillip Gibbons, Srikanta Tirthapura
HIGHLIGHT: We present an algorithm STRSAGA for efficiently maintaining a machine learning model over data points that arrive over time, quickly updating the model as new training data is observed.

911, TITLE: Online Robust Policy Learning in the Presence of Unknown Adversaries
https://papers.nips.cc/paper/8197-online-robust-policy-learning-in-the-presence-of-unknown-adversaries
AUTHORS: Aaron Havens, Zhanhong Jiang, Soumik Sarkar
HIGHLIGHT: This paper introduces a Meta-Learned Advantage Hierarchy (MLAH) framework that is attack model-agnostic and more suited to reinforcement learning, via handling the attacks in the decision space (as opposed to data space) and directly mitigating learned bias introduced by the adversary.

912, TITLE: Uplift Modeling from Separate Labels
https://papers.nips.cc/paper/8198-uplift-modeling-from-separate-labels
AUTHORS: Ikko Yamane, Florian Yger, Jamal Atif, Masashi Sugiyama
HIGHLIGHT: In this paper, we propose a novel method of uplift modeling that is applicable to a more practical setting where only one type of labels is available for each instance.

913, TITLE: Learning Invariances using the Marginal Likelihood
https://papers.nips.cc/paper/8199-learning-invariances-using-the-marginal-likelihood
AUTHORS: Mark van der Wilk, Matthias Bauer, ST John, James Hensman
HIGHLIGHT: We incorporate invariances in a Gaussian process, due to goodmarginal likelihood approximations being available for these models.

914, TITLE: Non-delusional Q-learning and value-iteration
https://papers.nips.cc/paper/8200-non-delusional-q-learning-and-value-iteration
AUTHORS: Tyler Lu, Dale Schuurmans, Craig Boutilier
HIGHLIGHT: To solve this problem, we introduce a new notion of policy consistency and define a local backup process that ensures global consistency through the use of information sets---sets that record constraints on policies consistent with backed-up Q-values.

915, TITLE: Using Large Ensembles of Control Variates for Variational Inference
https://papers.nips.cc/paper/8201-using-large-ensembles-of-control-variates-for-variational-inference
AUTHORS: Tomas Geffner, Justin Domke
HIGHLIGHT: In this paper we clarify the large number of control variates that are available by giving a systematic view of how they are derived.

916, TITLE: Post: Device Placement with Cross-Entropy Minimization and Proximal Policy Optimization
https://papers.nips.cc/paper/8202-post-device-placement-with-cross-entropy-minimization-and-proximal-policy-optimization
AUTHORS: Yuanxiang Gao, Li Chen, Baochun Li
HIGHLIGHT: In this paper, we propose a new joint learning algorithm, called Post, that integrates cross-entropy minimization and proximal policy optimization to achieve theoretically guaranteed optimal efficiency.

917, TITLE: Learning to Reason with Third Order Tensor Products
https://papers.nips.cc/paper/8203-learning-to-reason-with-third-order-tensor-products
AUTHORS: Imanol Schlag, J?rgen Schmidhuber
HIGHLIGHT: We combine Recurrent Neural Networks with Tensor Product Representations to learn combinatorial representations of sequential data.

918, TITLE: Memory Augmented Policy Optimization for Program Synthesis and Semantic Parsing
https://papers.nips.cc/paper/8204-memory-augmented-policy-optimization-for-program-synthesis-and-semantic-parsing
AUTHORS: Chen Liang, Mohammad Norouzi, Jonathan Berant, Quoc V. Le, Ni Lao
HIGHLIGHT: We present Memory Augmented Policy Optimization (MAPO), a simple and novel way to leverage a memory buffer of promising trajectories to reduce the variance of policy gradient estimate.

919, TITLE: Persistence Fisher Kernel: A Riemannian Manifold Kernel for Persistence Diagrams
https://papers.nips.cc/paper/8205-persistence-fisher-kernel-a-riemannian-manifold-kernel-for-persistence-diagrams
AUTHORS: Tam Le, Makoto Yamada
HIGHLIGHT: In this work, we rely upon the alternative \textit{Fisher information geometry} to propose a positive definite kernel for PDs \textit{without approximation}, namely the Persistence Fisher (PF) kernel.

920, TITLE: Neural Voice Cloning with a Few Samples
https://papers.nips.cc/paper/8206-neural-voice-cloning-with-a-few-samples
AUTHORS: Sercan Arik, Jitong Chen, Kainan Peng, Wei Ping, Yanqi Zhou
HIGHLIGHT: We introduce a neural voice cloning system that learns to synthesize a person's voice from only a few audio samples.

921, TITLE: Blind Deconvolutional Phase Retrieval via Convex Programming
https://papers.nips.cc/paper/8207-blind-deconvolutional-phase-retrieval-via-convex-programming
AUTHORS: Ali Ahmed, Alireza Aghasi, Paul Hand
HIGHLIGHT: We consider the task of recovering two real or complex m-vectors from phaseless Fourier measurements of their circular convolution.

922, TITLE: Scalable Laplacian K-modes
https://papers.nips.cc/paper/8208-scalable-laplacian-k-modes
AUTHORS: Imtiaz Ziko, Eric Granger, Ismail Ben Ayed
HIGHLIGHT: We advocate Laplacian K-modes for joint clustering and density mode finding, and propose a concave-convex relaxation of the problem, which yields a parallel algorithm that scales up to large datasets and high dimensions.

923, TITLE: A Retrieve-and-Edit Framework for Predicting Structured Outputs
https://papers.nips.cc/paper/8209-a-retrieve-and-edit-framework-for-predicting-structured-outputs
AUTHORS: Tatsunori B. Hashimoto, Kelvin Guu, Yonatan Oren, Percy S. Liang
HIGHLIGHT: With this motivation, we propose an approach that first retrieves a training example based on the input (e.g., natural language description) and then edits it to the desired output (e.g., code).

924, TITLE: Testing for Families of Distributions via the Fourier Transform
https://papers.nips.cc/paper/8210-testing-for-families-of-distributions-via-the-fourier-transform
AUTHORS: Cl?ment L. Canonne, Ilias Diakonikolas, Alistair Stewart
HIGHLIGHT: The main contribution of this work is a simple and general testing technique that is applicable to all distribution families whose Fourier spectrum satisfies a certain approximate sparsity property.

925, TITLE: Thwarting Adversarial Examples: An <var>L_0</var>-Robust Sparse Fourier Transform
https://papers.nips.cc/paper/8211-thwarting-adversarial-examples-an-l_0-robust-sparse-fourier-transform
AUTHORS: Mitali Bafna, Jack Murtagh, Nikhil Vyas
HIGHLIGHT: We give a new algorithm for approximating the Discrete Fourier transform of an approximately sparse signal that is robust to worst-case L_0 corruptions, namely that some coordinates of the signal can be corrupt arbitrarily.

926, TITLE: Blockwise Parallel Decoding for Deep Autoregressive Models
https://papers.nips.cc/paper/8212-blockwise-parallel-decoding-for-deep-autoregressive-models
AUTHORS: Mitchell Stern, Noam Shazeer, Jakob Uszkoreit
HIGHLIGHT: To overcome this limitation, we propose a novel blockwise parallel decoding scheme in which we make predictions for multiple time steps in parallel then back off to the longest prefix validated by a scoring model.

927, TITLE: Low-Rank Tucker Decomposition of Large Tensors Using TensorSketch
https://papers.nips.cc/paper/8213-low-rank-tucker-decomposition-of-large-tensors-using-tensorsketch
AUTHORS: Osman Asif Malik, Stephen Becker
HIGHLIGHT: We propose two randomized algorithms for low-rank Tucker decomposition of tensors.

928, TITLE: A Simple Cache Model for Image Recognition
https://papers.nips.cc/paper/8214-a-simple-cache-model-for-image-recognition
AUTHORS: Emin Orhan
HIGHLIGHT: We propose to extract this extra class-relevant information using a simple key-value cache memory to improve the classification performance of the model at test time.

929, TITLE: Clebsch?Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network
https://papers.nips.cc/paper/8215-clebschgordan-nets-a-fully-fourier-space-spherical-convolutional-neural-network
AUTHORS: Risi Kondor, Zhen Lin, Shubhendu Trivedi
HIGHLIGHT: In this paper we propose a generalization of this work that generally exhibits improved performace, but from an implementation point of view is actually simpler.

930, TITLE: Bayesian Nonparametric Spectral Estimation
https://papers.nips.cc/paper/8216-bayesian-nonparametric-spectral-estimation
AUTHORS: Felipe Tobar
HIGHLIGHT: In this context, we propose a joint probabilistic model for signals, observations and spectra, where SE is addressed as an inference problem.

931, TITLE: A Spectral View of Adversarially Robust Features
https://papers.nips.cc/paper/8217-a-spectral-view-of-adversarially-robust-features
AUTHORS: Shivam Garg, Vatsal Sharan, Brian Zhang, Gregory Valiant
HIGHLIGHT: Given the apparent difficulty of learning models that are robust to adversarial perturbations, we propose tackling the simpler problem of developing adversarially robust features.

932, TITLE: Synaptic Strength For Convolutional Neural Network
https://papers.nips.cc/paper/8218-synaptic-strength-for-convolutional-neural-network
AUTHORS: CHEN LIN, Zhao Zhong, Wu Wei, Junjie Yan
HIGHLIGHT: Our contribution is summarized as following: (1) We introduceSynaptic Strength, a new class of parameters for CNNs to indicate the importanceof each connections.

933, TITLE: Human-in-the-Loop Interpretability Prior
https://papers.nips.cc/paper/8219-human-in-the-loop-interpretability-prior
AUTHORS: Isaac Lage, Andrew Ross, Samuel J. Gershman, Been Kim, Finale Doshi-Velez
HIGHLIGHT: In this work, we optimize for interpretability by directly including humans in the optimization loop.

934, TITLE: Learning To Learn Around A Common Mean
https://papers.nips.cc/paper/8220-learning-to-learn-around-a-common-mean
AUTHORS: Giulia Denevi, Carlo Ciliberto, Dimitris Stamos, Massimiliano Pontil
HIGHLIGHT: In this work, we consider the family of algorithms given by a variant of Ridge Regression, in which the regularizer is the square distance to an unknown mean vector.

935, TITLE: Backpropagation with Callbacks: Foundations for Efficient and Expressive Differentiable Programming
https://papers.nips.cc/paper/8221-backpropagation-with-callbacks-foundations-for-efficient-and-expressive-differentiable-programming
AUTHORS: Fei Wang, James Decker, Xilun Wu, Gregory Essertel, Tiark Rompf
HIGHLIGHT: In this paper we propose an implementation of backpropagation using functions with callbacks, where the forward pass is executed as a sequence of function calls, and the backward pass as a corresponding sequence of function returns.

936, TITLE: Learning with SGD and Random Features
https://papers.nips.cc/paper/8222-learning-with-sgd-and-random-features
AUTHORS: Luigi Carratino, Alessandro Rudi, Lorenzo Rosasco
HIGHLIGHT: In this paper, we investigate their application in the context of nonparametric statistical learning.

937, TITLE: Total stochastic gradient algorithms and applications in reinforcement learning
https://papers.nips.cc/paper/8223-total-stochastic-gradient-algorithms-and-applications-in-reinforcement-learning
AUTHORS: Paavo Parmas
HIGHLIGHT: In this work we show how the total derivative rule leads to an intuitive visual framework for creating gradient estimators on graphical models.

938, TITLE: Glow: Generative Flow with Invertible 1x1 Convolutions
https://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions
AUTHORS: Durk P. Kingma, Prafulla Dhariwal
HIGHLIGHT: In this paper we propose Glow, a simple type of generative flow using invertible 1x1 convolution.

939, TITLE: Nonparametric Density Estimation under Adversarial Losses
https://papers.nips.cc/paper/8225-nonparametric-density-estimation-under-adversarial-losses
AUTHORS: Shashank Singh, Ananya Uppal, Boyue Li, Chun-Liang Li, Manzil Zaheer, Barnabas Poczos
HIGHLIGHT: In a general framework, we study how the choice of loss and the assumed smoothness of the underlying density together determine the minimax rate.

940, TITLE: Generalizing Point Embeddings using the Wasserstein Space of Elliptical Distributions

https://papers.nips.cc/paper/8226-generalizing-point-embeddings-using-the-wasserstein-space-of-elliptical-distributions
AUTHORS: Boris Muzellec, Marco Cuturi
HIGHLIGHT: We propose in this work an extension of that approach, which consists in embedding objects as elliptical probability distributions, namely distributions whose densities have elliptical level sets.

941, TITLE: Learning to Share and Hide Intentions using Information Regularization
https://papers.nips.cc/paper/8227-learning-to-share-and-hide-intentions-using-information-regularization
AUTHORS: DJ Strouse, Max Kleiman-Weiner, Josh Tenenbaum, Matt Botvinick, David J. Schwab
HIGHLIGHT: Here we show how to learn effective strategies for cooperation and competition in an asymmetric information game with no such model or interaction.

942, TITLE: Predictive Approximate Bayesian Computation via Saddle Points
https://papers.nips.cc/paper/8228-predictive-approximate-bayesian-computation-via-saddle-points
AUTHORS: Yingxiang Yang, Bo Dai, Negar Kiyavash, Niao He
HIGHLIGHT: In this paper, we introduce an optimization-based ABC framework that addresses these deficiencies.

943, TITLE: Robustness of conditional GANs to noisy labels
https://papers.nips.cc/paper/8229-robustness-of-conditional-gans-to-noisy-labels
AUTHORS: Kiran K. Thekumparampil, Ashish Khetan, Zinan Lin, Sewoong Oh
HIGHLIGHT: When the distribution of the noise is known, we introduce a novel architecture which we call Robust Conditional GAN (RCGAN).

944, TITLE: Robust Learning of Fixed-Structure Bayesian Networks
https://papers.nips.cc/paper/8230-robust-learning-of-fixed-structure-bayesian-networks
AUTHORS: Yu Cheng, Ilias Diakonikolas, Daniel Kane, Alistair Stewart
HIGHLIGHT: In this work, we study the fully observable discrete case where the structure of the network is given.

945, TITLE: Improving Simple Models with Confidence Profiles
https://papers.nips.cc/paper/8231-improving-simple-models-with-confidence-profiles
AUTHORS: Amit Dhurandhar, Karthikeyan Shanmugam, Ronny Luss, Peder A. Olsen
HIGHLIGHT: In this paper, we propose a new method called ProfWeight for transferring information from a pre-trained deep neural network that has a high test accuracy to a simpler interpretable model or a very shallow network of low complexity and a priori low test accuracy.

946, TITLE: PCA of high dimensional random walks with comparison to neural network training
https://papers.nips.cc/paper/8232-pca-of-high-dimensional-random-walks-with-comparison-to-neural-network-training
AUTHORS: Joseph Antognini, Jascha Sohl-Dickstein
HIGHLIGHT: In this paper we compare this technique to the PCA of a high dimensional random walk.

947, TITLE: Learning to Solve SMT Formulas
https://papers.nips.cc/paper/8233-learning-to-solve-smt-formulas
AUTHORS: Mislav Balunovic, Pavol Bielik, Martin Vechev
HIGHLIGHT: We present a new approach for learning to solve SMT formulas.

948, TITLE: Lifted Weighted Mini-Bucket
https://papers.nips.cc/paper/8234-lifted-weighted-mini-bucket
AUTHORS: Nicholas Gallo, Alexander T. Ihler
HIGHLIGHT: In this paper, we present a lifted variant of the Weighted Mini-Bucket elimination algorithm which provides a principled way to (i) exploit the highly symmetric substructure of MLN models, and (ii) incorporate high-order inference terms which are necessary for high quality approximate inference.

949, TITLE: Learning and Inference in Hilbert Space with Quantum Graphical Models
https://papers.nips.cc/paper/8235-learning-and-inference-in-hilbert-space-with-quantum-graphical-models
AUTHORS: Siddarth Srinivasan, Carlton Downey, Byron Boots
HIGHLIGHT: We show that these operations can be kernelized, and use these insights to propose a Hilbert Space Embedding of Hidden Quantum Markov Models (HSE-HQMM) to model dynamics.

950, TITLE: Unsupervised Image-to-Image Translation Using Domain-Specific Variational Information Bound
https://papers.nips.cc/paper/8236-unsupervised-image-to-image-translation-using-domain-specific-variational-information-bound
AUTHORS: Hadi Kazemi, Sobhan Soleymani, Fariborz Taherkhani, Seyed Iranmanesh, Nasser Nasrabadi

HIGHLIGHT: In this work, we propose an unsupervised image-to-image translation framework which maximizes a domain-specific variational information bound and learns the target domain-invariant representation of the two domain.

951, TITLE: Adversarial Risk and Robustness: General Definitions and Implications for the Uniform Distribution
https://papers.nips.cc/paper/8237-adversarial-risk-and-robustness-general-definitions-and-implications-for-the-uniform-distribution
AUTHORS: Dimitrios Diochnos, Saeed Mahloujifar, Mohammad Mahmoody
HIGHLIGHT: We study adversarial perturbations when the instances are uniformly distributed over {0,1}^n.

952, TITLE: Gaussian Process Prior Variational Autoencoders
https://papers.nips.cc/paper/8238-gaussian-process-prior-variational-autoencoders
AUTHORS: Francesco Paolo Casale, Adrian Dalca, Luca Saglietti, Jennifer Listgarten, Nicolo Fusi
HIGHLIGHT: In this work, we introduce a new model, the Gaussian Process (GP) Prior Variational Autoencoder (GPPVAE), to specifically address this issue.

953, TITLE: 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data
https://papers.nips.cc/paper/8239-3d-steerable-cnns-learning-rotationally-equivariant-features-in-volumetric-data
AUTHORS: Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, Taco Cohen
HIGHLIGHT: We present a convolutional network that is equivariant to rigid body motions.

954, TITLE: Context-aware Synthesis and Placement of Object Instances
https://papers.nips.cc/paper/8240-context-aware-synthesis-and-placement-of-object-instances
AUTHORS: Donghoon Lee, Sifei Liu, Jinwei Gu, Ming-Yu Liu, Ming-Hsuan Yang, Jan Kautz
HIGHLIGHT: In this paper, we propose an end-to-end trainable neural network for the task of inserting an object instance mask of a specified class into the semantic label map of an image.

955, TITLE: Convex Elicitation of Continuous Properties
https://papers.nips.cc/paper/8241-convex-elicitation-of-continuous-properties
AUTHORS: Jessica Finocchiaro, Rafael Frongillo
HIGHLIGHT: In this paper, in a finite-outcome setting, we show that in fact every elicitable real-valued property can be elicited by a convex loss function.

956, TITLE: Mesh-TensorFlow: Deep Learning for Supercomputers
https://papers.nips.cc/paper/8242-mesh-tensorflow-deep-learning-for-supercomputers
AUTHORS: Noam Shazeer, Youlong Cheng, Niki Parmar, Dustin Tran, Ashish Vaswani, Penporn Koanantakool, Peter Hawkins, HyoukJoong Lee, Mingsheng Hong, Cliff Young, Ryan Sepassi, Blake Hechtman
HIGHLIGHT: We introduce Mesh-TensorFlow, a language for specifying a general class of distributed tensor computations.

957, TITLE: Learning Abstract Options
https://papers.nips.cc/paper/8243-learning-abstract-options
AUTHORS: Matthew Riemer, Miao Liu, Gerald Tesauro
HIGHLIGHT: In this work, we extend previous work on this topic that only focuses on learning a two-level hierarchy including options and primitive actions to enable learning simultaneously at multiple resolutions in time.

958, TITLE: Bounded-Loss Private Prediction Markets
https://papers.nips.cc/paper/8244-bounded-loss-private-prediction-markets
AUTHORS: Rafael Frongillo, Bo Waggoner
HIGHLIGHT: In this work, we design an adaptively-growing prediction market with a bounded financial subsidy, while achieving privacy, incentives to produce accurate predictions, and precision in the sense that market prices are not heavily impacted by the added privacy-preserving noise.

959, TITLE: Temporal alignment and latent Gaussian process factor inference in population spike trains
https://papers.nips.cc/paper/8245-temporal-alignment-and-latent-gaussian-process-factor-inference-in-population-spike-trains
AUTHORS: Lea Duncker, Maneesh Sahani
HIGHLIGHT: We introduce a novel scalable approach to identifying common latent structure in neural population spike-trains, which allows for variability both in the trajectory and in the rate of progression of the underlying computation.

960, TITLE: Using Trusted Data to Train Deep Networks on Labels Corrupted by Severe Noise
https://papers.nips.cc/paper/8246-using-trusted-data-to-train-deep-networks-on-labels-corrupted-by-severe-noise
AUTHORS: Dan Hendrycks, Mantas Mazeika, Duncan Wilson, Kevin Gimpel

HIGHLIGHT: We demonstrate that robustness to label noise up to severe strengths can be achieved by using a set of trusted data with clean labels, and propose a loss correction that utilizes trusted examples in a data-efficient manner to mitigate the effects of label noise on deep neural network classifiers.

961, TITLE: Discretely Relaxing Continuous Variables for tractable Variational Inference
https://papers.nips.cc/paper/8247-discretely-relaxing-continuous-variables-for-tractable-variational-inference
AUTHORS: Trefor Evans, Prasanth Nair
HIGHLIGHT: We explore a new research direction in Bayesian variational inference with discrete latent variable priors where we exploit Kronecker matrix algebra for efficient and exact computations of the evidence lower bound (ELBO).

962, TITLE: Regret bounds for meta Bayesian optimization with an unknown Gaussian process prior
https://papers.nips.cc/paper/8248-regret-bounds-for-meta-bayesian-optimization-with-an-unknown-gaussian-process-prior
AUTHORS: Zi Wang, Beomjoon Kim, Leslie Pack Kaelbling
HIGHLIGHT: In this paper, we adopt a variant of empirical Bayes and show that, by estimating the Gaussian process prior from offline data sampled from the same prior and constructing unbiased estimators of the posterior, variants of both GP-UCB and \emph{probability of improvement} achieve a near-zero regret bound, which decreases to a constant proportional to the observational noise as the number of offline data and the number of online evaluations increase.

963, TITLE: Diversity-Driven Exploration Strategy for Deep Reinforcement Learning
https://papers.nips.cc/paper/8249-diversity-driven-exploration-strategy-for-deep-reinforcement-learning
AUTHORS: Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, Tsu-Jui Fu, Chun-Yi Lee
HIGHLIGHT: To tackle this problem, we present a diversity-driven approach for exploration, which can be easily combined with both off- and on-policy reinforcement learning algorithms.

964, TITLE: Deep Generative Models with Learnable Knowledge Constraints
https://papers.nips.cc/paper/8250-deep-generative-models-with-learnable-knowledge-constraints
AUTHORS: Zhiting Hu, Zichao Yang, Ruslan R. Salakhutdinov, LIANHUI Qin, Xiaodan Liang, Haoye Dong, Eric P. Xing
HIGHLIGHT: In this paper, we establish mathematical correspondence between PR and reinforcement learning (RL), and, based on the connection, expand PR to learn constraints as the extrinsic reward in RL.

965, TITLE: The Sparse Manifold Transform
https://papers.nips.cc/paper/8251-the-sparse-manifold-transform
AUTHORS: Yubei Chen, Dylan Paiton, Bruno Olshausen
HIGHLIGHT: We present a signal representation framework called the sparse manifold transform that combines key ideas from sparse coding, manifold learning, and slow feature analysis.

966, TITLE: Bayesian Structure Learning by Recursive Bootstrap
https://papers.nips.cc/paper/8252-bayesian-structure-learning-by-recursive-bootstrap
AUTHORS: Raanan Y. Rohekar, Yaniv Gurwicz, Shami Nisimov, Guy Koren, Gal Novik
HIGHLIGHT: We propose a method that covers both model averaging and model selection in the same framework.

967, TITLE: Complex Gated Recurrent Neural Networks
https://papers.nips.cc/paper/8253-complex-gated-recurrent-neural-networks
AUTHORS: Moritz Wolter, Angela Yao
HIGHLIGHT: We present a novel complex gated recurrent cell, which is a hybrid cell combining complex-valued and norm-preserving state transitions with a gating mechanism.

968, TITLE: Learning a Warping Distance from Unlabeled Time Series Using Sequence Autoencoders
https://papers.nips.cc/paper/8254-learning-a-warping-distance-from-unlabeled-time-series-using-sequence-autoencoders
AUTHORS: Abubakar Abid, James Y. Zou
HIGHLIGHT: In this paper, we propose an end-to-end framework, autowarp, that optimizes and learns a good metric given unlabeled trajectories.

969, TITLE: Streamlining Variational Inference for Constraint Satisfaction Problems
https://papers.nips.cc/paper/8255-streamlining-variational-inference-for-constraint-satisfaction-problems
AUTHORS: Aditya Grover, Tudor Achim, Stefano Ermon
HIGHLIGHT: We introduce a more general branching strategy based on streamlining constraints, which sidestep hard assignments to variables.

970, TITLE: Fast deep reinforcement learning using online adjustments from the past

https://papers.nips.cc/paper/8256-fast-deep-reinforcement-learning-using-online-adjustments-from-the-past
AUTHORS: Steven Hansen, Alexander Pritzel, Pablo Sprechmann, Andre Barreto, Charles Blundell
HIGHLIGHT: We propose Ephemeral Value Adjusments (EVA): a means of allowing deep reinforcement learning agents to rapidly adapt to experience in their replay buffer.

971, TITLE: Improved Network Robustness with Adversary Critic
https://papers.nips.cc/paper/8257-improved-network-robustness-with-adversary-critic
AUTHORS: Alexander Matyasko, Lap-Pui Chau
HIGHLIGHT: To address this gap in perception, we propose a novel approach for learning robust classifier.

972, TITLE: Regret Bounds for Online Portfolio Selection with a Cardinality Constraint
https://papers.nips.cc/paper/8258-regret-bounds-for-online-portfolio-selection-with-a-cardinality-constraint
AUTHORS: Shinji Ito, Daisuke Hatano, Sumita Hanna, Akihiro Yabe, Takuro Fukunaga, Naonori Kakimura, Ken-Ichi Kawarabayashi
HIGHLIGHT: In this paper, we study the problem with the cardinality constraint that the number of assets in a portfolio is restricted to be at most k, and consider two scenarios: (i) in the full-feedback setting, the learner can observe price relatives (rates of return to cost) for all assets, and (ii) in the bandit-feedback setting, the learner can observe price relatives only for invested assets.

973, TITLE: Sketching Method for Large Scale Combinatorial Inference
https://papers.nips.cc/paper/8259-sketching-method-for-large-scale-combinatorial-inference
AUTHORS: Wei Sun, Junwei Lu, Han Liu
HIGHLIGHT: We present computationally efficient algorithms to test various combinatorial structures of large-scale graphical models.

974, TITLE: Connecting Optimization and Regularization Paths
https://papers.nips.cc/paper/8260-connecting-optimization-and-regularization-paths
AUTHORS: Arun Suggala, Adarsh Prasad, Pradeep K. Ravikumar
HIGHLIGHT: We study the implicit regularization properties of optimization techniques by explicitly connecting their optimization paths to the regularization paths of ``corresponding'' regularized problems.

975, TITLE: Fully Neural Network Based Speech Recognition on Mobile and Embedded Devices
https://papers.nips.cc/paper/8261-fully-neural-network-based-speech-recognition-on-mobile-and-embedded-devices
AUTHORS: Jinhwan Park, Yoonho Boo, Iksoo Choi, Sungho Shin, Wonyong Sung
HIGHLIGHT: We present real-time speech recognition on smartphones or embedded systems by employing recurrent neural network (RNN) based acoustic models, RNN based language models, and beam-search decoding.

976, TITLE: Understanding Regularized Spectral Clustering via Graph Conductance
https://papers.nips.cc/paper/8262-understanding-regularized-spectral-clustering-via-graph-conductance
AUTHORS: Yilin Zhang, Karl Rohe
HIGHLIGHT: This paper uses the relationship between graph conductance and spectral clustering to study (i) the failures of spectral clustering and (ii) the benefits of regularization.

977, TITLE: Data-Driven Clustering via Parameterized Lloyd's Families
https://papers.nips.cc/paper/8263-data-driven-clustering-via-parameterized-lloyds-families
AUTHORS: Maria-Florina F. Balcan, Travis Dick, Colin White
HIGHLIGHT: In this paper, we define an infinite family of algorithms generalizing Lloyd's algorithm, with one parameter controlling the the initialization procedure, and another parameter controlling the local search procedure.

978, TITLE: Learning Beam Search Policies via Imitation Learning
https://papers.nips.cc/paper/8264-learning-beam-search-policies-via-imitation-learning
AUTHORS: Renato Negrinho, Matthew Gormley, Geoffrey J. Gordon
HIGHLIGHT: We develop an unifying meta-algorithm for learning beam search policies using imitation learning.

979, TITLE: Benefits of over-parameterization with EM
https://papers.nips.cc/paper/8265-benefits-of-over-parameterization-with-em
AUTHORS: Ji Xu, Daniel J. Hsu, Arian Maleki
HIGHLIGHT: The goal of this article is to present theoretical and empirical evidence that over-parameterization can help EM avoid spurious local optima in the log-likelihood.

980, TITLE: Thermostat-assisted continuously-tempered Hamiltonian Monte Carlo for Bayesian learning

https://papers.nips.cc/paper/8266-thermostat-assisted-continuously-tempered-hamiltonian-monte-carlo-for-bayesian-learning
AUTHORS: Rui Luo, Jianhong Wang, Yaodong Yang, Jun WANG, Zhanxing Zhu
HIGHLIGHT: In this paper, we propose a novel sampling method, the thermostat-assisted continuously-tempered Hamiltonian Monte Carlo, for the purpose of multimodal Bayesian learning.

981, TITLE: Robust Subspace Approximation in a Stream
https://papers.nips.cc/paper/8267-robust-subspace-approximation-in-a-stream
AUTHORS: Roie Levin, Anish Prasad Sevekari, David Woodruff
HIGHLIGHT: We give the first sublinear approximation algorithm for this problem in the turnstile streaming and arbitrary partition distributed models, achieving the same time guarantees as in the offline case.

982, TITLE: Mean Field for the Stochastic Blockmodel: Optimization Landscape and Convergence Issues
https://papers.nips.cc/paper/8268-mean-field-for-the-stochastic-blockmodel-optimization-landscape-and-convergence-issues
AUTHORS: Soumendu Sundar Mukherjee, Purnamrita Sarkar, Y. X. Rachel Wang, Bowei Yan
HIGHLIGHT: In this paper, we focus on the problem of community detection for a simple two-class Stochastic Blockmodel (SBM).

983, TITLE: Analysis of Krylov Subspace Solutions of Regularized Non-Convex Quadratic Problems
https://papers.nips.cc/paper/8269-analysis-of-krylov-subspace-solutions-of-regularized-non-convex-quadratic-problems
AUTHORS: Yair Carmon, John C. Duchi
HIGHLIGHT: We provide convergence rates for Krylov subspace solutions to the trust-region and cubic-regularized (nonconvex) quadratic problems.

984, TITLE: Autoconj: Recognizing and Exploiting Conjugacy Without a Domain-Specific Language
https://papers.nips.cc/paper/8270-autoconj-recognizing-and-exploiting-conjugacy-without-a-domain-specific-language
AUTHORS: Matthew D. Hoffman, Matthew J. Johnson, Dustin Tran
HIGHLIGHT: In this paper, we propose a strategy for automating such derivations.

985, TITLE: DropBlock: A regularization method for convolutional networks
https://papers.nips.cc/paper/8271-dropblock-a-regularization-method-for-convolutional-networks
AUTHORS: Golnaz Ghiasi, Tsung-Yi Lin, Quoc V. Le
HIGHLIGHT: In this paper, we introduce DropBlock, a form of structured dropout, where units in a contiguous region of a feature map are dropped together.

986, TITLE: Forward Modeling for Partial Observation Strategy Games - A StarCraft Defogger
https://papers.nips.cc/paper/8272-forward-modeling-for-partial-observation-strategy-games-a-starcraft-defogger
AUTHORS: Gabriel Synnaeve, Zeming Lin, Jonas Gehring, Dan Gant, Vegard Mella, Vasil Khalidov, Nicolas Carion, Nicolas Usunier
HIGHLIGHT: We propose to employ encoder-decoder neural networks for this task, and introduce proxy tasks and baselines for evaluation to assess their ability of capturing basic game rules and high-level dynamics.

987, TITLE: With Friends Like These, Who Needs Adversaries?
https://papers.nips.cc/paper/8273-with-friends-like-these-who-needs-adversaries
AUTHORS: Saumya Jetley, Nicholas Lord, Philip Torr
HIGHLIGHT: This provides a novel perspective on the existence of universal adversarial perturbations.

988, TITLE: Decentralize and Randomize: Faster Algorithm for Wasserstein Barycenters
https://papers.nips.cc/paper/8274-decentralize-and-randomize-faster-algorithm-for-wasserstein-barycenters
AUTHORS: Pavel Dvurechenskii, Darina Dvinskikh, Alexander Gasnikov, Cesar Uribe, Angelia Nedich
HIGHLIGHT: Motivated by this problem, we develop, and analyze, a novel accelerated primal-dual stochastic gradient method for general stochastic convex optimization problems with linear equality constraints.

989, TITLE: Joint Autoregressive and Hierarchical Priors for Learned Image Compression
https://papers.nips.cc/paper/8275-joint-autoregressive-and-hierarchical-priors-for-learned-image-compression
AUTHORS: David Minnen, Johannes Ball?, George D. Toderici
HIGHLIGHT: Recently, hierarchical entropy models were introduced as a way to exploit more structure in the latents than previous fully factorized priors, improving compression performance while maintaining end-to-end optimization.

990, TITLE: Learning Temporal Point Processes via Reinforcement Learning
https://papers.nips.cc/paper/8276-learning-temporal-point-processes-via-reinforcement-learning

AUTHORS: Shuang Li, Shuai Xiao, Shixiang Zhu, Nan Du, Yao Xie, Le Song
HIGHLIGHT: To alleviate the risk of model-misspecification in MLE, we propose to generate samples from the generative model and monitor the quality of the samples in the process of training until the samples and the real data are indistinguishable.

991, TITLE: Bias and Generalization in Deep Generative Models: An Empirical Study
https://papers.nips.cc/paper/8277-bias-and-generalization-in-deep-generative-models-an-empirical-study
AUTHORS: Shengjia Zhao, Hongyu Ren, Arianna Yuan, Jiaming Song, Noah Goodman, Stefano Ermon
HIGHLIGHT: In this paper we propose a framework to systematically investigate bias and generalization in deep generative models of images by probing the learning algorithm with carefully designed training datasets.

992, TITLE: Fast and Effective Robustness Certification
https://papers.nips.cc/paper/8278-fast-and-effective-robustness-certification
AUTHORS: Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus P?schel, Martin Vechev
HIGHLIGHT: We present a new method and system, called DeepZ, for certifying neural network robustness based on abstract interpretation.

993, TITLE: Support Recovery for Orthogonal Matching Pursuit: Upper and Lower bounds
https://papers.nips.cc/paper/8279-support-recovery-for-orthogonal-matching-pursuit-upper-and-lower-bounds
AUTHORS: Raghav Somani, Chirag Gupta, Prateek Jain, Praneeth Netrapalli
HIGHLIGHT: This paper studies the problem of sparse regression where the goal is to learn a sparse vector that best optimizes a given objective function.

994, TITLE: Differentially Private Change-Point Detection
https://papers.nips.cc/paper/8280-differentially-private-change-point-detection
AUTHORS: Rachel Cummings, Sara Krehbiel, Yajun Mei, Rui Tuo, Wanrong Zhang
HIGHLIGHT: We study the statistical problem of change-point problem through the lens of differential privacy.

995, TITLE: Multi-value Rule Sets for Interpretable Classification with Feature-Efficient Representations
https://papers.nips.cc/paper/8281-multi-value-rule-sets-for-interpretable-classification-with-feature-efficient-representations
AUTHORS: Tong Wang
HIGHLIGHT: We propose a Bayesian framework for formulating an MRS model and develop an efficient inference method for learning a maximum a posteriori, incorporating theoretically grounded bounds to iteratively reduce the search space and improve the search efficiency.

996, TITLE: Domain Adaptation by Using Causal Inference to Predict Invariant Conditional Distributions
https://papers.nips.cc/paper/8282-domain-adaptation-by-using-causal-inference-to-predict-invariant-conditional-distributions
AUTHORS: Sara Magliacane, Thijs van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, Joris M. Mooij
HIGHLIGHT: We propose an approach for solving these problems that exploits causal inference and does not rely on prior knowledge of the causal graph, the type of interventions or the intervention targets.

997, TITLE: Smoothed Analysis of Discrete Tensor Decomposition and Assemblies of Neurons
https://papers.nips.cc/paper/8283-smoothed-analysis-of-discrete-tensor-decomposition-and-assemblies-of-neurons
AUTHORS: Nima Anari, Constantinos Daskalakis, Wolfgang Maass, Christos Papadimitriou, Amin Saberi, Santosh Vempala
HIGHLIGHT: We analyze linear independence of rank one tensors produced by tensor powers of randomly perturbed vectors.

998, TITLE: MixLasso: Generalized Mixed Regression via Convex Atomic-Norm Regularization
https://papers.nips.cc/paper/8284-mixlasso-generalized-mixed-regression-via-convex-atomic-norm-regularization
AUTHORS: Ian En-Hsu Yen, Wei-Cheng Lee, Kai Zhong, Sung-En Chang, Pradeep K. Ravikumar, Shou-De Lin
HIGHLIGHT: In this work, we study a novel convex estimator \emph{MixLasso} for the estimation of generalized mixed regression, based on an atomic norm specifically constructed to regularize the number of mixture components.

999, TITLE: Semidefinite relaxations for certifying robustness to adversarial examples
https://papers.nips.cc/paper/8285-semidefinite-relaxations-for-certifying-robustness-to-adversarial-examples
AUTHORS: Aditi Raghunathan, Jacob Steinhardt, Percy S. Liang
HIGHLIGHT: In this paper, we propose a new semidefinite relaxation for certifying robustness that applies to arbitrary ReLU networks.

1000, TITLE: Removing Hidden Confounding by Experimental Grounding
https://papers.nips.cc/paper/8286-removing-hidden-confounding-by-experimental-grounding

AUTHORS: Nathan Kallus, Aahlad Manas Puli, Uri Shalit
HIGHLIGHT: We introduce a novel method of using limited experimental data to correct the hidden confounding in causal effect models trained on larger observational data, even if the observational data does not fully overlap with the experimental data.

1001, TITLE: Topkapi: Parallel and Fast Sketches for Finding Top-K Frequent Elements
https://papers.nips.cc/paper/8287-topkapi-parallel-and-fast-sketches-for-finding-top-k-frequent-elements
AUTHORS: Ankush Mandal, He Jiang, Anshumali Shrivastava, Vivek Sarkar
HIGHLIGHT: In this paper, we identify that in modern distributed settings with both multi-node as well as multi-core parallelism, existing algorithms, although theoretically sound, are suboptimal from the performance perspective.

1002, TITLE: Contrastive Learning from Pairwise Measurements
https://papers.nips.cc/paper/8288-contrastive-learning-from-pairwise-measurements
AUTHORS: Yi Chen, Zhuoran Yang, Yuchen Xie, Princeton Zhaoran Wang
HIGHLIGHT: In this paper, we study a semiparametric model where the pairwise measurements follow a natural exponential family distribution with an unknown base measure.

1003, TITLE: Point process latent variable models of larval zebrafish behavior
https://papers.nips.cc/paper/8289-point-process-latent-variable-models-of-larval-zebrafish-behavior
AUTHORS: Anuj Sharma, Robert Johnson, Florian Engert, Scott Linderman
HIGHLIGHT: We develop a new class of probabilistic models to tackle this challenge in the study of larval zebrafish, an important model organism for neuroscience.

1004, TITLE: Computationally and statistically efficient learning of causal Bayes nets using path queries
https://papers.nips.cc/paper/8290-computationally-and-statistically-efficient-learning-of-causal-bayes-nets-using-path-queries
AUTHORS: Kevin Bello, Jean Honorio
HIGHLIGHT: We theoretically show the logarithmic sample complexity for the size of interventional data per path query, for continuous and discrete networks.

1005, TITLE: Sparse PCA from Sparse Linear Regression
https://papers.nips.cc/paper/8291-sparse-pca-from-sparse-linear-regression
AUTHORS: Guy Bresler, Sung Min Park, Madalina Persu
HIGHLIGHT: We show how to efficiently transform a black-box solver for SLR into an algorithm for SPCA: assuming the SLR solver satisfies prediction error guarantees achieved by existing efficient algorithms such as those based on the Lasso, the SPCA algorithm derived from it achieves near state of the art guarantees for testing and for support recovery for the single spiked covariance model as obtained by the current best polynomial-time algorithms.

1006, TITLE: Multiple Instance Learning for Efficient Sequential Data Classification on Resource-constrained Devices
https://papers.nips.cc/paper/8292-multiple-instance-learning-for-efficient-sequential-data-classification-on-resource-constrained-devices
AUTHORS: Don Dennis, Chirag Pabbaraju, Harsha Vardhan Simhadri, Prateek Jain
HIGHLIGHT: In this paper, we address this challenge by exploiting the following two observations about classification tasks arising in typical IoT related applications: (a) the "signature" of a particular class (e.g. an audio keyword) typically occupies a small fraction of the overall data, and (b) class signatures tend to be discernible early on in the data.

1007, TITLE: Transfer of Deep Reactive Policies for MDP Planning
https://papers.nips.cc/paper/8293-transfer-of-deep-reactive-policies-for-mdp-planning
AUTHORS: Aniket (Nick) Bajpai, Sankalp Garg, Mausam
HIGHLIGHT: In this paper, we present the first domain-independent transfer algorithm for MDP planning domains expressed in an RDDL representation.

1008, TITLE: The Price of Fair PCA: One Extra dimension
https://papers.nips.cc/paper/8294-the-price-of-fair-pca-one-extra-dimension
AUTHORS: Samira Samadi, Uthaipon Tantipongpipat, Jamie H. Morgenstern, Mohit Singh, Santosh Vempala
HIGHLIGHT: We define the notion of Fair PCA and give a polynomial-time algorithm for finding a low dimensional representation of the data which is nearly-optimal with respect to this measure.

1009, TITLE: GroupReduce: Block-Wise Low-Rank Approximation for Neural Language Model Shrinking
https://papers.nips.cc/paper/8295-groupreduce-block-wise-low-rank-approximation-for-neural-language-model-shrinking
AUTHORS: Patrick Chen, Si Si, Yang Li, Ciprian Chelba, Cho-Jui Hsieh

HIGHLIGHT: In this paper, we propose GroupReduce, a novel compression method for neural language models, based on vocabulary-partition (block) based low-rank matrix approximation and the inherent frequency distribution of tokens (the power-law distribution of words).