

- 1, TITLE: Wider and Deeper, Cheaper and Faster: Tensorized LSTMs for Sequence Learning
<https://papers.nips.cc/paper/6606-wider-and-deeper-cheaper-and-faster-tensorized-lstms-for-sequence-learning>
AUTHORS: Zhen He, Shaobing Gao, Liang Xiao, Daxue Liu, Hangen He, David Barber
HIGHLIGHT: As an alternative we propose the Tensorized LSTM in which the hidden states are represented by tensors and updated via a cross-layer convolution.
- 2, TITLE: Concentration of Multilinear Functions of the Ising Model with Applications to Network Data
<https://papers.nips.cc/paper/6607-concentration-of-multilinear-functions-of-the-ising-model-with-applications-to-network-data>
AUTHORS: Constantinos Daskalakis, Nishanth Dikkala, Gautam Kamath
HIGHLIGHT: We prove near-tight concentration of measure for polynomial functions of the Ising model, under high temperature, improving the radius of concentration guaranteed by known results by polynomial factors in the dimension (i.e.~the number of nodes in the Ising model).
- 3, TITLE: Deep Subspace Clustering Networks
<https://papers.nips.cc/paper/6608-deep-subspace-clustering-networks>
AUTHORS: Pan Ji, Tong Zhang, Hongdong Li, Mathieu Salzmann, Ian Reid
HIGHLIGHT: We present a novel deep neural network architecture for unsupervised subspace clustering.
- 4, TITLE: Attentional Pooling for Action Recognition
<https://papers.nips.cc/paper/6609-attentional-pooling-for-action-recognition>
AUTHORS: Rohit Girdhar, Deva Ramanan
HIGHLIGHT: We introduce a simple yet surprisingly powerful model to incorporate attention in action recognition and human object interaction tasks.
- 5, TITLE: On the Consistency of Quick Shift
<https://papers.nips.cc/paper/6610-on-the-consistency-of-quick-shift>
AUTHORS: Heinrich Jiang
HIGHLIGHT: We present finite sample statistical consistency guarantees for Quick Shift on mode and cluster recovery under mild distributional assumptions.
- 6, TITLE: Breaking the Nonsmooth Barrier: A Scalable Parallel Method for Composite Optimization
<https://papers.nips.cc/paper/6611-breaking-the-nonsmooth-barrier-a-scalable-parallel-method-for-composite-optimization>
AUTHORS: Fabian Pedregosa, R?mi Leblond, Simon Lacoste-Julien
HIGHLIGHT: In this work, we propose and analyze ProxASAGA, a fully asynchronous sparse method inspired by SAGA, a variance reduced incremental gradient algorithm.
- 7, TITLE: Dual-Agent GANs for Photorealistic and Identity Preserving Profile Face Synthesis
<https://papers.nips.cc/paper/6612-dual-agent-gans-for-photorealistic-and-identity-preserving-profile-face-synthesis>
AUTHORS: Jian Zhao, Lin Xiong, Panasonic Karlekar Jayashree, Jianshu Li, Fang Zhao, Zhecan Wang, Panasonic Sugiri Pranata, Panasonic Shengmei Shen, Shuicheng Yan, Jiashi Feng
HIGHLIGHT: To narrow this gap, we propose a Dual-Agent Generative Adversarial Network (DA-GAN) model, which can improve the realism of a face simulator's output using unlabeled real faces, while preserving the identity information during the realism refinement.
- 8, TITLE: Dilated Recurrent Neural Networks
<https://papers.nips.cc/paper/6613-dilated-recurrent-neural-networks>
AUTHORS: Shiyu Chang, Yang Zhang, Wei Han, Mo Yu, Xiaoxiao Guo, Wei Tan, Xiaodong Cui, Michael Witbrock, Mark A. Hasegawa-Johnson, Thomas S. Huang
HIGHLIGHT: In this paper, we introduce a simple yet effective RNN connection structure, the DilatedRNN, which simultaneously tackles all of these challenges.
- 9, TITLE: Hunt For The Unique, Stable, Sparse And Fast Feature Learning On Graphs
<https://papers.nips.cc/paper/6614-hunt-for-the-unique-stable-sparse-and-fast-feature-learning-on-graphs>
AUTHORS: Saurabh Verma, Zhi-Li Zhang
HIGHLIGHT: To both evaluate the quality of graph features produced by FGSD and demonstrate their utility, we apply them to the graph classification problem.
- 10, TITLE: Scalable Generalized Linear Bandits: Online Computation and Hashing
<https://papers.nips.cc/paper/6615-scalable-generalized-linear-bandits-online-computation-and-hashing>
AUTHORS: Kwang-Sung Jun, Aniruddha Bhargava, Robert Nowak, Rebecca Willett
HIGHLIGHT: This paper proposes new, scalable solutions to the GLB problem in two respects.

- 11, TITLE: Probabilistic Models for Integration Error in the Assessment of Functional Cardiac Models
<https://papers.nips.cc/paper/6616-probabilistic-models-for-integration-error-in-the-assessment-of-functional-cardiac-models>
AUTHORS: Chris Oates, Steven Niederer, Angela Lee, Francois-Xavier Briol, Mark Girolami
HIGHLIGHT: This paper studies the numerical computation of integrals, representing estimates or predictions, over the output $f(x)$ of a computational model with respect to a distribution $p(\mathbf{d}x)$ over uncertain inputs x to the model.
- 12, TITLE: Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent
<https://papers.nips.cc/paper/6617-machine-learning-with-adversaries-byzantine-tolerant-gradient-descent>
AUTHORS: Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, Julien Stainer
HIGHLIGHT: We propose Krum , an aggregation rule that satisfies our resilience property, which we argue is the first provably Byzantine-resilient algorithm for distributed SGD.
- 13, TITLE: Dynamic Safe Interruptibility for Decentralized Multi-Agent Reinforcement Learning
<https://papers.nips.cc/paper/6618-dynamic-safe-interruptibility-for-decentralized-multi-agent-reinforcement-learning>
AUTHORS: El Mahdi El Mhamdi, Rachid Guerraoui, Hadrien Hendrikx, Alexandre Maurer
HIGHLIGHT: This paper introduces dynamic safe interruptibility, an alternative definition more suited to decentralized learning problems, and studies this notion in two learning frameworks: joint action learners and independent learners.
- 14, TITLE: Interactive Submodular Bandit
<https://papers.nips.cc/paper/6619-interactive-submodular-bandit>
AUTHORS: Lin Chen, Andreas Krause, Amin Karbasi
HIGHLIGHT: We model such problems as an interactive submodular bandit optimization, where in each round we receive a context (e.g., previously selected movies) and have to choose an action (e.g., propose a new movie).
- 15, TITLE: Learning to See Physics via Visual De-animation
<https://papers.nips.cc/paper/6620-learning-to-see-physics-via-visual-de-animation>
AUTHORS: Jiajun Wu, Erika Lu, Pushmeet Kohli, Bill Freeman, Josh Tenenbaum
HIGHLIGHT: We introduce a paradigm for understanding physical scenes without human annotations.
- 16, TITLE: Label Efficient Learning of Transferable Representations across Domains and Tasks
<https://papers.nips.cc/paper/6621-label-efficient-learning-of-transferable-representations-across-domains-and-tasks>
AUTHORS: Zelun Luo, Yuliang Zou, Judy Hoffman, Li F. Fei-Fei
HIGHLIGHT: We propose a framework that learns a representation transferable across different domains and tasks in a data efficient manner.
- 17, TITLE: Decoding with Value Networks for Neural Machine Translation
<https://papers.nips.cc/paper/6622-decoding-with-value-networks-for-neural-machine-translation>
AUTHORS: Di He, Hanqing Lu, Yingce Xia, Tao Qin, Liwei Wang, Tie-Yan Liu
HIGHLIGHT: Inspired by the success and methodology of AlphaGo, in this paper we propose using a prediction network to improve beam search, which takes the source sentence x , the currently available decoding output y_1, \dots, y_{t-1} and a candidate word w at step t as inputs and predicts the long-term value (e.g., BLEU score) of the partial target sentence if it is completed by the NMT model.
- 18, TITLE: Parametric Simplex Method for Sparse Learning
<https://papers.nips.cc/paper/6623-parametric-simplex-method-for-sparse-learning>
AUTHORS: Haotian Pang, Han Liu, Robert J. Vanderbei, Tuo Zhao
HIGHLIGHT: In this paper, we investigate a broad class of sparse learning approaches formulated as linear programs parametrized by a regularization factor, and solve them by the parametric simplex method (PSM).
- 19, TITLE: Group Sparse Additive Machine
<https://papers.nips.cc/paper/6624-group-sparse-additive-machine>
AUTHORS: Hong Chen, Xiaoqian Wang, Cheng Deng, Heng Huang
HIGHLIGHT: To address this challenging problem, in this paper, we investigate the classification with group sparse additive models in reproducing kernel Hilbert spaces.
- 20, TITLE: Uprooting and Rerooting Higher-Order Graphical Models
<https://papers.nips.cc/paper/6625-uprooting-and-rerooting-higher-order-graphical-models>
AUTHORS: Mark Rowland, Adrian Weller

- HIGHLIGHT:** Here we introduce methods to extend the approach to models with higher-order potentials and develop theoretical insights.
- 21, **TITLE:** The Unreasonable Effectiveness of Structured Random Orthogonal Embeddings
<https://papers.nips.cc/paper/6626-the-unreasonable-effectiveness-of-structured-random-orthogonal-embeddings>
AUTHORS: Krzysztof M. Choromanski, Mark Rowland, Adrian Weller
HIGHLIGHT: We introduce matrices with complex entries which give significant further accuracy improvement.
- 22, **TITLE:** From Parity to Preference-based Notions of Fairness in Classification
<https://papers.nips.cc/paper/6627-from-parity-to-preference-based-notions-of-fairness-in-classification>
AUTHORS: Muhammad Bilal Zafar, Isabel Valera, Manuel Rodriguez, Krishna Gummadi, Adrian Weller
HIGHLIGHT: In this paper, we draw inspiration from the fair-division and envy-freeness literature in economics and game theory and propose preference-based notions of fairness -- given the choice between various sets of decision treatments or outcomes, any group of users would collectively prefer its treatment or outcomes, regardless of the (dis)parity as compared to the other groups.
- 23, **TITLE:** Inferring Generative Model Structure with Static Analysis
<https://papers.nips.cc/paper/6628-inferring-generative-model-structure-with-static-analysis>
AUTHORS: Paroma Varma, Bryan D. He, Payal Bajaj, Nishith Khandwala, Imon Banerjee, Daniel Rubin, Christopher R?
HIGHLIGHT: We present Coral, a paradigm that infers generative model structure by statically analyzing the code for these heuristics, thus significantly reducing the amount of data required to learn structure.
- 24, **TITLE:** Structured Embedding Models for Grouped Data
<https://papers.nips.cc/paper/6629-structured-embedding-models-for-grouped-data>
AUTHORS: Maja Rudolph, Francisco Ruiz, Susan Athey, David Blei
HIGHLIGHT: Here we develop structured exponential family embeddings (S-EFE), a method for discovering embeddings that vary across related groups of data.
- 25, **TITLE:** A Linear-Time Kernel Goodness-of-Fit Test
<https://papers.nips.cc/paper/6630-a-linear-time-kernel-goodness-of-fit-test>
AUTHORS: Wittawat Jitkrittum, Wenkai Xu, Zoltan Szabo, Kenji Fukumizu, Arthur Gretton
HIGHLIGHT: We propose a novel adaptive test of goodness-of-fit, with computational cost linear in the number of samples.
- 26, **TITLE:** Cortical microcircuits as gated-recurrent neural networks
<https://papers.nips.cc/paper/6631-cortical-microcircuits-as-gated-recurrent-neural-networks>
AUTHORS: Rui Costa, Ioannis Alexandros Assael, Brendan Shillingford, Nando de Freitas, Tim Vogels
HIGHLIGHT: We propose a natural mapping of subLSTMs onto known canonical excitatory-inhibitory cortical microcircuits.
- 27, **TITLE:** k-Support and Ordered Weighted Sparsity for Overlapping Groups: Hardness and Algorithms
<https://papers.nips.cc/paper/6632-k-support-and-ordered-weighted-sparsity-for-overlapping-groups-hardness-and-algorithms>
AUTHORS: Cong Han Lim, Stephen Wright
HIGHLIGHT: We study the norms obtained from extending the k-support norm and OWL norms to the setting in which there are overlapping groups.
- 28, **TITLE:** A simple model of recognition and recall memory
<https://papers.nips.cc/paper/6633-a-simple-model-of-recognition-and-recall-memory>
AUTHORS: Nisheeth Srivastava, Edward Vul
HIGHLIGHT: We build a simple computational model around this theory, using sampling to approximate an ideal Bayesian observer encoding and retrieving situational co-occurrence frequencies of stimuli and retrieval cues.
- 29, **TITLE:** On Structured Prediction Theory with Calibrated Convex Surrogate Losses
<https://papers.nips.cc/paper/6634-on-structured-prediction-theory-with-calibrated-convex-surrogate-losses>
AUTHORS: Anton Osokin, Francis Bach, Simon Lacoste-Julien
HIGHLIGHT: We provide novel theoretical insights on structured prediction in the context of efficient convex surrogate loss minimization with consistency guarantees.
- 30, **TITLE:** Best of Both Worlds: Transferring Knowledge from Discriminative Learning to a Generative Visual Dialog Model
<https://papers.nips.cc/paper/6635-best-of-both-worlds-transferring-knowledge-from-discriminative-learning-to-a-generative-visual-dialog-model>
AUTHORS: Jiasen Lu, Anitha Kannan, Jianwei Yang, Devi Parikh, Dhruv Batra

HIGHLIGHT: We present a novel training framework for neural sequence models, particularly for grounded dialog generation.

31, TITLE: MaskRNN: Instance Level Video Object Segmentation

<https://papers.nips.cc/paper/6636-maskrnn-instance-level-video-object-segmentation>

AUTHORS: Yuan-Ting Hu, Jia-Bin Huang, Alexander Schwing

HIGHLIGHT: To capture the temporal coherence, in this paper, we develop MaskRNN, a recurrent neural net approach which fuses in each frame the output of two deep nets for each object instance - a binary segmentation net providing a mask and a localization net providing a bounding box.

32, TITLE: Gated Recurrent Convolution Neural Network for OCR

<https://papers.nips.cc/paper/6637-gated-recurrent-convolution-neural-network-for-ocr>

AUTHORS: Jianfeng Wang, Xiaolin Hu

HIGHLIGHT: Inspired by a recently proposed model for general image classification, Recurrent Convolution Neural Network (RCNN), we propose a new architecture named Gated RCNN (GRCNN) for solving this problem.

33, TITLE: Towards Accurate Binary Convolutional Neural Network

<https://papers.nips.cc/paper/6638-towards-accurate-binary-convolutional-neural-network>

AUTHORS: Xiaofan Lin, Cong Zhao, Wei Pan

HIGHLIGHT: In this paper, we address this issue with two major innovations: (1) approximating full-precision weights with the linear combination of multiple binary weight bases; (2) employing multiple binary activations to alleviate information loss.

34, TITLE: Semi-Supervised Learning for Optical Flow with Generative Adversarial Networks

<https://papers.nips.cc/paper/6639-semi-supervised-learning-for-optical-flow-with-generative-adversarial-networks>

AUTHORS: Wei-Sheng Lai, Jia-Bin Huang, Ming-Hsuan Yang

HIGHLIGHT: In this paper, we propose to exploit unlabeled videos for semi-supervised learning of optical flow with a Generative Adversarial Network.

35, TITLE: Learning a Multi-View Stereo Machine

<https://papers.nips.cc/paper/6640-learning-a-multi-view-stereo-machine>

AUTHORS: Abhishek Kar, Christian H?ne, Jitendra Malik

HIGHLIGHT: We present a learnt system for multi-view stereopsis.

36, TITLE: Phase Transitions in the Pooled Data Problem

<https://papers.nips.cc/paper/6641-phase-transitions-in-the-pooled-data-problem>

AUTHORS: Jonathan Scarlett, Volkan Cevher

HIGHLIGHT: In this paper, we study the {em pooled data} problem of identifying the labels associated with a large collection of items, based on a sequence of pooled tests revealing the counts of each label within the pool.

37, TITLE: Universal Style Transfer via Feature Transforms

<https://papers.nips.cc/paper/6642-universal-style-transfer-via-feature-transforms>

AUTHORS: Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, Ming-Hsuan Yang

HIGHLIGHT: In this paper, we present a simple yet effective method that tackles these limitations without training on any pre-defined styles.

38, TITLE: On the Model Shrinkage Effect of Gamma Process Edge Partition Models

<https://papers.nips.cc/paper/6643-on-the-model-shrinkage-effect-of-gamma-process-edge-partition-models>

AUTHORS: Iku Ohama, Issei Sato, Takuya Kida, Hiroki Arimura

HIGHLIGHT: In order to ensure that the model shrinkage effect of the EPM works in an appropriate manner, we proposed two novel generative constructions of the EPM: CEPM incorporating constrained gamma priors, and DEPM incorporating Dirichlet priors instead of the gamma priors.

39, TITLE: Pose Guided Person Image Generation

<https://papers.nips.cc/paper/6644-pose-guided-person-image-generation>

AUTHORS: Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, Luc Van Gool

HIGHLIGHT: This paper proposes the novel Pose Guided Person Generation Network (PGS²) that allows to synthesize person images in arbitrary poses, based on an image of that person and a novel pose.

40, TITLE: Inference in Graphical Models via Semidefinite Programming Hierarchies

<https://papers.nips.cc/paper/6645-inference-in-graphical-models-via-semidefinite-programming-hierarchies>

AUTHORS: Murat A. Erdogdu, Yash Deshpande, Andrea Montanari

HIGHLIGHT: In this paper, we propose binary SDP relaxations for MAP inference using the SOS hierarchy with two innovations focused on computational efficiency.

41, **TITLE:** Variable Importance Using Decision Trees
<https://papers.nips.cc/paper/6646-variable-importance-using-decision-trees>
AUTHORS: Jalil Kazemitabar, Arash Amini, Adam Bloniarz, Ameet S. Talwalkar
HIGHLIGHT: Decision trees and random forests are well established models that not only offer good predictive performance, but also provide rich feature importance information.

42, **TITLE:** Preventing Gradient Explosions in Gated Recurrent Units
<https://papers.nips.cc/paper/6647-preventing-gradient-explosions-in-gated-recurrent-units>
AUTHORS: Sekitoshi Kanai, Yasuhiro Fujiwara, Sotetsu Iwamura
HIGHLIGHT: In this paper, we find a condition under which the dynamics of the GRU changes drastically and propose a learning method to address the exploding gradient problem.

43, **TITLE:** On the Power of Truncated SVD for General High-rank Matrix Estimation Problems
<https://papers.nips.cc/paper/6648-on-the-power-of-truncated-svd-for-general-high-rank-matrix-estimation-problems>
AUTHORS: Simon S. Du, Yining Wang, Aarti Singh
HIGHLIGHT: 2.High-rank matrix denoising: we design algorithms that recovers a matrix \mathbf{A} with relative error in Frobenius norm from its noise-perturbed observations, without assuming \mathbf{A} is exactly low-rank.

44, **TITLE:** f-GANs in an Information Geometric Nutshell
<https://papers.nips.cc/paper/6649-f-gans-in-an-information-geometric-nutshell>
AUTHORS: Richard Nock, Zac Cranko, Aditya K. Menon, Lizhen Qu, Robert C. Williamson
HIGHLIGHT: In this paper, we unveil a broad class of distributions for which such convergence happens --- namely, deformed exponential families, a wide superset of exponential families ---.

45, **TITLE:** Toward Multimodal Image-to-Image Translation
<https://papers.nips.cc/paper/6650-toward-multimodal-image-to-image-translation>
AUTHORS: Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A. Efros, Oliver Wang, Eli Shechtman
HIGHLIGHT: In this work, we aim to model a distribution of possible outputs in a conditional generative modeling setting.

46, **TITLE:** Mixture-Rank Matrix Approximation for Collaborative Filtering
<https://papers.nips.cc/paper/6651-mixture-rank-matrix-approximation-for-collaborative-filtering>
AUTHORS: Dongsheng Li, Chao Chen, Wei Liu, Tun Lu, Ning Gu, Stephen Chu
HIGHLIGHT: In this paper, a mixture-rank matrix approximation (MRMA) method is proposed, in which user-item ratings can be characterized by a mixture of LRMA models with different ranks.

47, **TITLE:** Continuous DR-submodular Maximization: Structure and Algorithms
<https://papers.nips.cc/paper/6652-continuous-dr-submodular-maximization-structure-and-algorithms>
AUTHORS: An Bian, Kfir Levy, Andreas Krause, Joachim M. Buhmann
HIGHLIGHT: In this work we study the problem of maximizing non-monotone DR-submodular continuous functions under general down-closed convex constraints.

48, **TITLE:** Learning with Average Top-k Loss
<https://papers.nips.cc/paper/6653-learning-with-average-top-k-loss>
AUTHORS: Yanbo Fan, Siwei Lyu, Yiming Ying, Baogang Hu
HIGHLIGHT: In this work, we introduce the average top- k (atk) loss as a new ensemble loss for supervised learning.

49, **TITLE:** Learning multiple visual domains with residual adapters
<https://papers.nips.cc/paper/6654-learning-multiple-visual-domains-with-residual-adapters>
AUTHORS: Sylvestre-Alvise Rebuffi, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: In this paper, we look in particular at the task of learning a single visual representation that can be successfully utilized in the analysis of very different types of images, from dog breeds to stop signs and digits. We also introduce the Visual Decathlon Challenge, a benchmark that evaluates the ability of representations to capture simultaneously ten very different visual domains and measures their ability to recognize well uniformly.

50, **TITLE:** Dykstra's Algorithm, ADMM, and Coordinate Descent: Connections, Insights, and Extensions
<https://papers.nips.cc/paper/6655-dykstras-algorithm-admm-and-coordinate-descent-connections-insights-and-extensions>
AUTHORS: Ryan J. Tibshirani

HIGHLIGHT: We study connections between Dykstra's algorithm for projecting onto an intersection of convex sets, the augmented Lagrangian method of multipliers or ADMM, and block coordinate descent.

51, **TITLE:** Learning Spherical Convolution for Fast Features from 360° Imagery
<https://papers.nips.cc/paper/6656-learning-spherical-convolution-for-fast-features-from-360-imagery>

AUTHORS: Yu-Chuan Su, Kristen Grauman

HIGHLIGHT: We propose to learn a spherical convolutional network that translates a planar CNN to process 360° imagery directly in its equirectangular projection.

52, **TITLE:** MarrNet: 3D Shape Reconstruction via 2.5D Sketches

<https://papers.nips.cc/paper/6657-marrnet-3d-shape-reconstruction-via-25d-sketches>

AUTHORS: Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, Josh Tenenbaum

HIGHLIGHT: In this work, we propose an end-to-end trainable framework, sequentially estimating 2.5D sketches and 3D object shapes.

53, **TITLE:** Multimodal Learning and Reasoning for Visual Question Answering

<https://papers.nips.cc/paper/6658-multimodal-learning-and-reasoning-for-visual-question-answering>

AUTHORS: Ilija Ilievski, Jiashi Feng

HIGHLIGHT: In this work we introduce a modular neural network model that learns a multimodal and multifaceted representation of the image and the question.

54, **TITLE:** Adversarial Surrogate Losses for Ordinal Regression

<https://papers.nips.cc/paper/6659-adversarial-surrogate-losses-for-ordinal-regression>

AUTHORS: Rizal Fathony, Mohammad Ali Bashiri, Brian Ziebart

HIGHLIGHT: Many existing methods for this task reduce to binary classification problems and employ surrogate losses, such as the hinge loss.

55, **TITLE:** Hypothesis Transfer Learning via Transformation Functions

<https://papers.nips.cc/paper/6660-hypothesis-transfer-learning-via-transformation-functions>

AUTHORS: Simon S. Du, Jayanth Koushik, Aarti Singh, Barnabas Poczos

HIGHLIGHT: In this paper, we propose a unified algorithm-dependent framework for HTL through a novel notion of transformation functions, which characterizes the relation between the source and the target domains.

56, **TITLE:** Controllable Invariance through Adversarial Feature Learning

<https://papers.nips.cc/paper/6661-controllable-invariance-through-adversarial-feature-learning>

AUTHORS: Qizhe Xie, Zihang Dai, Yulun Du, Eduard Hovy, Graham Neubig

HIGHLIGHT: In this paper, we tackle the problem of learning representations invariant to a specific factor or trait of data.

57, **TITLE:** Convergence Analysis of Two-layer Neural Networks with ReLU Activation

<https://papers.nips.cc/paper/6662-convergence-analysis-of-two-layer-neural-networks-with-relu-activation>

AUTHORS: Yuanzhi Li, Yang Yuan

HIGHLIGHT: In this paper, we make progress on understanding this mystery by providing a convergence analysis for SGD on a rich subset of two-layer feedforward networks with ReLU activations.

58, **TITLE:** Doubly Accelerated Stochastic Variance Reduced Dual Averaging Method for Regularized Empirical Risk Minimization

<https://papers.nips.cc/paper/6663-doubly-accelerated-stochastic-variance-reduced-dual-averaging-method-for-regularized-empirical-risk-minimization>

AUTHORS: Tomoya Murata, Taiji Suzuki

HIGHLIGHT: We develop a new accelerated stochastic gradient method for efficiently solving the convex regularized empirical risk minimization problem in mini-batch settings.

59, **TITLE:** Langevin Dynamics with Continuous Tempering for Training Deep Neural Networks

<https://papers.nips.cc/paper/6664-langevin-dynamics-with-continuous-tempering-for-training-deep-neural-networks>

AUTHORS: Nanyang Ye, Zhanxing Zhu, Rafal Mantiuk

HIGHLIGHT: In this paper, a novel approach is proposed which divides the training process into two consecutive phases to obtain better generalization performance: Bayesian sampling and stochastic optimization.

60, **TITLE:** Efficient Online Linear Optimization with Approximation Algorithms

<https://papers.nips.cc/paper/6665-efficient-online-linear-optimization-with-approximation-algorithms>

AUTHORS: Dan Garber
HIGHLIGHT: We present new algorithms with significantly improved oracle complexity for both the full information and bandit variants of the problem.

61, TITLE: Geometric Descent Method for Convex Composite Minimization
<https://papers.nips.cc/paper/6666-geometric-descent-method-for-convex-composite-minimization>
AUTHORS: Shixiang Chen, Shiqian Ma, Wei Liu
HIGHLIGHT: In this paper, we extend the geometric descent method recently proposed by Bubeck, Lee and Singh to tackle nonsmooth and strongly convex composite problems.

62, TITLE: Diffusion Approximations for Online Principal Component Estimation and Global Convergence
<https://papers.nips.cc/paper/6667-diffusion-approximations-for-online-principal-component-estimation-and-global-convergence>
AUTHORS: Chris Junchi Li, Mengdi Wang, Han Liu, Tong Zhang
HIGHLIGHT: In this paper, we propose to adopt the diffusion approximation tools to study the dynamics of Oja's iteration which is an online stochastic gradient method for the principal component analysis.

63, TITLE: Avoiding Discrimination through Causal Reasoning
<https://papers.nips.cc/paper/6668-avoiding-discrimination-through-causal-reasoning>
AUTHORS: Niki Kilbertus, Mateo Rojas Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, Bernhard Schölkopf
HIGHLIGHT: Going beyond observational criteria, we frame the problem of discrimination based on protected attributes in the language of causal reasoning.

64, TITLE: Nonparametric Online Regression while Learning the Metric
<https://papers.nips.cc/paper/6669-nonparametric-online-regression-while-learning-the-metric>
AUTHORS: Ilja Kuzborskij, Nicolò Cesa-Bianchi
HIGHLIGHT: We study algorithms for online nonparametric regression that learn the directions along which the regression function is smoother.

65, TITLE: Recycling Privileged Learning and Distribution Matching for Fairness
<https://papers.nips.cc/paper/6670-recycling-privileged-learning-and-distribution-matching-for-fairness>
AUTHORS: Novi Quadrianto, Viktoriia Sharmanska
HIGHLIGHT: To achieve our goal, we recycle two well-established machine learning techniques, privileged learning and distribution matching, and harmonize them for satisfying multi-faceted fairness definitions.

66, TITLE: Safe and Nested Subgame Solving for Imperfect-Information Games
<https://papers.nips.cc/paper/6671-safe-and-nested-subgame-solving-for-imperfect-information-games>
AUTHORS: Noam Brown, Tuomas Sandholm
HIGHLIGHT: We introduce subgame-solving techniques that outperform prior methods both in theory and practice.

67, TITLE: Unsupervised Image-to-Image Translation Networks
<https://papers.nips.cc/paper/6672-unsupervised-image-to-image-translation-networks>
AUTHORS: Ming-Yu Liu, Thomas Breuel, Jan Kautz
HIGHLIGHT: To address the problem, we make a shared-latent space assumption and propose an unsupervised image-to-image translation framework based on Coupled GANs.

68, TITLE: Coded Distributed Computing for Inverse Problems
<https://papers.nips.cc/paper/6673-coded-distributed-computing-for-inverse-problems>
AUTHORS: Yaoqing Yang, Pulkit Grover, Soumya Kar
HIGHLIGHT: In this paper, we utilize the emerging idea of "coded computation" to design a novel error-correcting-code inspired technique for solving linear inverse problems under specific iterative methods in a parallelized implementation affected by stragglers.

69, TITLE: A Screening Rule for ℓ_1 -Regularized Ising Model Estimation
<https://papers.nips.cc/paper/6674-a-screening-rule-for-l1-regularized-ising-model-estimation>
AUTHORS: Zhaobin Kuang, Sinong Geng, David Page
HIGHLIGHT: We discover a screening rule for ℓ_1 -regularized Ising model estimation.

70, TITLE: Improved Dynamic Regret for Non-degenerate Functions
<https://papers.nips.cc/paper/6675-improved-dynamic-regret-for-non-degenerate-functions>

AUTHORS: Lijun Zhang, Tianbao Yang, Jinfeng Yi, Jing Rong, Zhi-Hua Zhou
HIGHLIGHT: In this paper, we illustrate that the dynamic regret can be further improved by allowing the learner to query the gradient of the function multiple times, and meanwhile the strong convexity can be weakened to other non-degenerate conditions.

71, TITLE: Learning Efficient Object Detection Models with Knowledge Distillation
<https://papers.nips.cc/paper/6676-learning-efficient-object-detection-models-with-knowledge-distillation>
AUTHORS: Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, Manmohan Chandraker
HIGHLIGHT: In this work, we propose a new framework to learn compact and fast object detection networks with improved accuracy using knowledge distillation [20] and hint learning [34].

72, TITLE: One-Sided Unsupervised Domain Mapping
<https://papers.nips.cc/paper/6677-one-sided-unsupervised-domain-mapping>
AUTHORS: Sagie Benaim, Lior Wolf
HIGHLIGHT: In this work, we present a method of learning $G_{\{AB\}}$ without learning $G_{\{BA\}}$.

73, TITLE: Deep Mean-Shift Priors for Image Restoration
<https://papers.nips.cc/paper/6678-deep-mean-shift-priors-for-image-restoration>
AUTHORS: Siavash Arjomand Bigdeli, Matthias Zwicker, Paolo Favaro, Meiguang Jin
HIGHLIGHT: In this paper we introduce a natural image prior that directly represents a Gaussian-smoothed version of the natural image distribution.

74, TITLE: Greedy Algorithms for Cone Constrained Optimization with Convergence Guarantees
<https://papers.nips.cc/paper/6679-greedy-algorithms-for-cone-constrained-optimization-with-convergence-guarantees>
AUTHORS: Francesco Locatello, Michael Tschannen, Gunnar Raetsch, Martin Jaggi
HIGHLIGHT: In this paper, we consider the intermediate case of optimization over the convex cone, parametrized as the conic hull of a generic atom set, leading to the first principled definitions of non-negative MP algorithms for which we give explicit convergence rates and demonstrate excellent empirical performance.

75, TITLE: A New Theory for Matrix Completion
<https://papers.nips.cc/paper/6680-a-new-theory-for-matrix-completion>
AUTHORS: Guangcan Liu, Qingshan Liu, Xiaotong Yuan
HIGHLIGHT: To break through the limits of random sampling, this paper introduces a new hypothesis called $\{isomeric\}$ condition, which is provably weaker than the assumption of uniform sampling and arguably holds even when the missing data is placed irregularly.

76, TITLE: Robust Hypothesis Test for Nonlinear Effect with Gaussian Processes
<https://papers.nips.cc/paper/6681-robust-hypothesis-test-for-nonlinear-effect-with-gaussian-processes>
AUTHORS: Jeremiah Liu, Brent Coull
HIGHLIGHT: To demonstrate the utility of the proposed method, we apply our test to the problem of detecting nonlinear interaction between groups of continuous features.

77, TITLE: Lower bounds on the robustness to adversarial perturbations
<https://papers.nips.cc/paper/6682-lower-bounds-on-the-robustness-to-adversarial-perturbations>
AUTHORS: Jonathan Peck, Joris Roels, Bart Goossens, Yvan Saeys
HIGHLIGHT: In this work, we take steps towards a formal characterization of adversarial perturbations by deriving lower bounds on the magnitudes of perturbations necessary to change the classification of neural networks.

78, TITLE: Minimizing a Submodular Function from Samples
<https://papers.nips.cc/paper/6683-minimizing-a-submodular-function-from-samples>
AUTHORS: Eric Balkanski, Yaron Singer
HIGHLIGHT: In this paper we consider the problem of minimizing a submodular function from training data.

79, TITLE: Introspective Classification with Convolutional Nets
<https://papers.nips.cc/paper/6684-introspective-classification-with-convolutional-nets>
AUTHORS: Long Jin, Justin Lazarow, Zhuowen Tu
HIGHLIGHT: We propose introspective convolutional networks (ICN) that emphasize the importance of having convolutional neural networks empowered with generative capabilities.

80, TITLE: Label Distribution Learning Forests
<https://papers.nips.cc/paper/6685-label-distribution-learning-forests>

AUTHORS: Wei Shen, KAI ZHAO, Yilu Guo, Alan L. Yuille
HIGHLIGHT: This paper presents label distribution learning forests (LDFs) - a novel label distribution learning algorithm based on differentiable decision trees, which have several advantages: 1) Decision trees have the potential to model any general form of label distributions by a mixture of leaf node predictions.

81, TITLE: Unsupervised learning of object frames by dense equivariant image labelling
<https://papers.nips.cc/paper/6686-unsupervised-learning-of-object-frames-by-dense-equivariant-image-labelling>
AUTHORS: James Thewlis, Hakan Bilen, Andrea Vedaldi
HIGHLIGHT: Starting from the recent idea of viewpoint factorization, we propose a new approach that, given a large number of images of an object and no other supervision, can extract a dense object-centric coordinate frame.

82, TITLE: Compression-aware Training of Deep Networks
<https://papers.nips.cc/paper/6687-compression-aware-training-of-deep-networks>
AUTHORS: Jose M. Alvarez, Mathieu Salzmann
HIGHLIGHT: In this paper, we propose to explicitly account for compression in the training process.

83, TITLE: Multiscale Semi-Markov Dynamics for Intracortical Brain-Computer Interfaces
<https://papers.nips.cc/paper/6688-multiscale-semi-markov-dynamics-for-intracortical-brain-computer-interfaces>
AUTHORS: Daniel Milstein, Jason Pacheco, Leigh Hochberg, John D. Simeral, Beata Jarosiewicz, Erik Sudderth
HIGHLIGHT: We propose a dynamic Bayesian network that includes the on-screen goal position as part of its latent state, and thus allows the person's intended angle of movement to be aggregated over a much longer history of neural activity.

84, TITLE: PredRNN: Recurrent Neural Networks for Predictive Learning using Spatiotemporal LSTMs
<https://papers.nips.cc/paper/6689-predrnn-recurrent-neural-networks-for-predictive-learning-using-spatiotemporal-lstms>
AUTHORS: Yunbo Wang, Mingsheng Long, Jianmin Wang, Zhifeng Gao, Philip S. Yu
HIGHLIGHT: PredRNN: Recurrent Neural Networks for Predictive Learning using Spatiotemporal LSTMs

85, TITLE: Detrended Partial Cross Correlation for Brain Connectivity Analysis
<https://papers.nips.cc/paper/6690-detrended-partial-cross-correlation-for-brain-connectivity-analysis>
AUTHORS: Jaime Ide, F?bio Cappabianco, Fabio Faria, Chiang-shan R. Li
HIGHLIGHT: In this paper, using detrended partial cross-correlation analysis (DPCCA), we propose a novel functional connectivity measure to delineate brain interactions at multiple time scales, while controlling for covariates.

86, TITLE: Contrastive Learning for Image Captioning
<https://papers.nips.cc/paper/6691-contrastive-learning-for-image-captioning>
AUTHORS: Bo Dai, Dahua Lin
HIGHLIGHT: In this work, we propose a new learning method, Contrastive Learning (CL), for image captioning.

87, TITLE: Safe Model-based Reinforcement Learning with Stability Guarantees
<https://papers.nips.cc/paper/6692-safe-model-based-reinforcement-learning-with-stability-guarantees>
AUTHORS: Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, Andreas Krause
HIGHLIGHT: In this paper, we present a learning algorithm that explicitly considers safety, defined in terms of stability guarantees.

88, TITLE: Online multiclass boosting
<https://papers.nips.cc/paper/6693-online-multiclass-boosting>
AUTHORS: Young Hun Jung, Jack Goetz, Ambuj Tewari
HIGHLIGHT: Recent work has extended the theoretical analysis of boosting algorithms to multiclass problems and to online settings.

89, TITLE: Matching on Balanced Nonlinear Representations for Treatment Effects Estimation
<https://papers.nips.cc/paper/6694-matching-on-balanced-nonlinear-representations-for-treatment-effects-estimation>
AUTHORS: Sheng Li, Yun Fu
HIGHLIGHT: In this paper, we aim to address these problems by learning low-dimensional balanced and nonlinear representations (BNR) for observational data.

90, TITLE: Learning Overcomplete HMMs
<https://papers.nips.cc/paper/6695-learning-overcomplete-hmms>
AUTHORS: Vatsal Sharan, Sham M. Kakade, Percy S. Liang, Gregory Valiant

HIGHLIGHT: In this paper, we present several new results---both positive and negative---which help define the boundaries between the tractable-learning setting and the intractable setting.

91, **TITLE:** GP CaKe: Effective brain connectivity with causal kernels
<https://papers.nips.cc/paper/6696-gp-cake-effective-brain-connectivity-with-causal-kernels>
AUTHORS: Luca Ambrogioni, Max Hinne, Marcel Van Gerven, Eric Maris
HIGHLIGHT: Here we propose to model this causal interaction using integro-differential equations and causal kernels that allow for a rich analysis of effective connectivity.

92, **TITLE:** Decoupling "when to update" from "how to update"
<https://papers.nips.cc/paper/6697-decoupling-when-to-update-from-how-to-update>
AUTHORS: Eran Malach, Shai Shalev-Shwartz
HIGHLIGHT: In this paper, we propose a meta algorithm for tackling the noisy labels problem.

93, **TITLE:** Self-Normalizing Neural Networks
<https://papers.nips.cc/paper/6698-self-normalizing-neural-networks>
AUTHORS: G?nter Klambauer, Thomas Unterthiner, Andreas Mayr, Sepp Hochreiter
HIGHLIGHT: We introduce self-normalizing neural networks (SNNs) to enable high-level abstract representations.

94, **TITLE:** Learning to Pivot with Adversarial Networks
<https://papers.nips.cc/paper/6699-learning-to-pivot-with-adversarial-networks>
AUTHORS: Gilles Louppe, Michael Kagan, Kyle Cranmer
HIGHLIGHT: In this work, we introduce and derive theoretical results for a training procedure based on adversarial networks for enforcing the pivotal property (or, equivalently, fairness with respect to continuous attributes) on a predictive model.

95, **TITLE:** SchNet: A continuous-filter convolutional neural network for modeling quantum interactions
<https://papers.nips.cc/paper/6700-schnet-a-continuous-filter-convolutional-neural-network-for-modeling-quantum-interactions>
AUTHORS: Kristof Sch?tt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre Tkatchenko, Klaus-Robert M?ller
HIGHLIGHT: Thus, we propose to use continuous-filter convolutional layers to be able to model local correlations without requiring the data to lie on a grid. Finally, we introduce a more challenging benchmark with chemical and structural variations that suggests the path for further work.

96, **TITLE:** Active Bias: Training More Accurate Neural Networks by Emphasizing High Variance Samples
<https://papers.nips.cc/paper/6701-active-bias-training-more-accurate-neural-networks-by-emphasizing-high-variance-samples>
AUTHORS: Haw-Shiuan Chang, Erik Learned-Miller, Andrew McCallum
HIGHLIGHT: This paper presents two improved alternatives based on lightweight estimates of sample uncertainty in stochastic gradient descent (SGD): the variance in predicted probability of the correct class across iterations of mini-batch SGD, and the proximity of the correct class probability to the decision threshold.

97, **TITLE:** Differentiable Learning of Submodular Models
<https://papers.nips.cc/paper/6702-differentiable-learning-of-submodular-models>
AUTHORS: Josip Djolonga, Andreas Krause
HIGHLIGHT: In this paper we focus on the problem of submodular minimization, for which we show that such layers are indeed possible.

98, **TITLE:** Inductive Representation Learning on Large Graphs
<https://papers.nips.cc/paper/6703-inductive-representation-learning-on-large-graphs>
AUTHORS: Will Hamilton, Zitao Ying, Jure Leskovec
HIGHLIGHT: Here we present GraphSAGE, a general, inductive framework that leverages node feature information (e.g., text attributes) to efficiently generate node embeddings.

99, **TITLE:** Subset Selection and Summarization in Sequential Data
<https://papers.nips.cc/paper/6704-subset-selection-and-summarization-in-sequential-data>
AUTHORS: Ehsan Elhamifar, M. Clara De Paolis Kaluza
HIGHLIGHT: In this paper, we develop a new framework for sequential subset selection that finds a set of representatives compatible with the dynamic models of data. To do so, we equip items with transition dynamic models and pose the problem as an integer binary optimization over assignments of sequential items to representatives, that leads to high encoding, diversity and transition potentials.

- 100, TITLE: Question Asking as Program Generation
<https://papers.nips.cc/paper/6705-question-asking-as-program-generation>
AUTHORS: Anselm Rothe, Brenden M. Lake, Todd Gureckis
HIGHLIGHT: Here we introduce a cognitive model capable of constructing human-like questions.
- 101, TITLE: Revisiting Perceptron: Efficient and Label-Optimal Learning of Halfspaces
<https://papers.nips.cc/paper/6706-revisiting-perceptron-efficient-and-label-optimal-learning-of-halfspaces>
AUTHORS: Songbai Yan, Chicheng Zhang
HIGHLIGHT: In this work, we propose an efficient Perceptron-based algorithm for actively learning homogeneous halfspaces under the uniform distribution over the unit sphere.
- 102, TITLE: Gradient Descent Can Take Exponential Time to Escape Saddle Points
<https://papers.nips.cc/paper/6707-gradient-descent-can-take-exponential-time-to-escape-saddle-points>
AUTHORS: Simon S. Du, Chi Jin, Jason D. Lee, Michael I. Jordan, Aarti Singh, Barnabas Poczos
HIGHLIGHT: Gradient Descent Can Take Exponential Time to Escape Saddle Points
- 103, TITLE: Union of Intersections (UoI) for Interpretable Data Driven Discovery and Prediction
<https://papers.nips.cc/paper/6708-union-of-intersections-uoI-for-interpretable-data-driven-discovery-and-prediction>
AUTHORS: Kristofer Bouchard, Alejandro Bujan, Farbod Roosta-Khorasani, Shashanka Ubaru, Mr. Prabhat, Antoine Snijders, Jian-Hua Mao, Edward Chang, Michael W. Mahoney, Sharmodeep Bhattacharya
HIGHLIGHT: We introduce the Union of Intersections (UoI) method, a flexible, modular, and scalable framework for enhanced model selection and estimation.
- 104, TITLE: One-Shot Imitation Learning
<https://papers.nips.cc/paper/6709-one-shot-imitation-learning>
AUTHORS: Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, Wojciech Zaremba
HIGHLIGHT: In this paper, we propose a meta-learning framework for achieving such capability, which we call one-shot imitation learning.
- 105, TITLE: Learning the Morphology of Brain Signals Using Alpha-Stable Convolutional Sparse Coding
<https://papers.nips.cc/paper/6710-learning-the-morphology-of-brain-signals-using-alpha-stable-convolutional-sparse-coding>
AUTHORS: Mainak Jas, Tom Dupr? la Tour, Umut Simsekli, Alexandre Gramfort
HIGHLIGHT: In this study, we address these issues and propose a novel probabilistic convolutional sparse coding (CSC) model for learning shift-invariant atoms from raw neural signals containing potentially severe artifacts.
- 106, TITLE: Integration Methods and Optimization Algorithms
<https://papers.nips.cc/paper/6711-integration-methods-and-optimization-algorithms>
AUTHORS: Damien Scieur, Vincent Roulet, Francis Bach, Alexandre d'Aspremont
HIGHLIGHT: We show that accelerated optimization methods can be seen as particular instances of multi-step integration schemes from numerical analysis, applied to the gradient flow equation.
- 107, TITLE: Sharpness, Restart and Acceleration
<https://papers.nips.cc/paper/6712-sharpness-restart-and-acceleration>
AUTHORS: Vincent Roulet, Alexandre d'Aspremont
HIGHLIGHT: The constants quantifying error bounds are of course unobservable, but we show that optimal restart strategies are robust, and searching for the best scheme only increases the complexity by a logarithmic factor compared to the optimal bound.
- 108, TITLE: Learning Koopman Invariant Subspaces for Dynamic Mode Decomposition
<https://papers.nips.cc/paper/6713-learning-koopman-invariant-subspaces-for-dynamic-mode-decomposition>
AUTHORS: Naoya Takeishi, Yoshinobu Kawahara, Takehisa Yairi
HIGHLIGHT: In this paper, we propose a fully data-driven method for Koopman spectral analysis based on the principle of learning Koopman invariant subspaces from observed data.
- 109, TITLE: Soft-to-Hard Vector Quantization for End-to-End Learning Compressible Representations
<https://papers.nips.cc/paper/6714-soft-to-hard-vector-quantization-for-end-to-end-learning-compressible-representations>
AUTHORS: Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, Luc V. Gool
HIGHLIGHT: We present a new approach to learn compressible representations in deep architectures with an end-to-end training strategy.

- 110, TITLE: Learning spatiotemporal piecewise-geodesic trajectories from longitudinal manifold-valued data
<https://papers.nips.cc/paper/6715-learning-spatiotemporal-piecewise-geodesic-trajectories-from-longitudinal-manifold-valued-data>
AUTHORS: St?phanie ALLASSONNIERE, Juliette Chevallier, Stephane Oudard
HIGHLIGHT: We introduce a hierarchical model which allows to estimate a group-average piecewise-geodesic trajectory in the Riemannian space of measurements and individual variability.
- 111, TITLE: Improving Regret Bounds for Combinatorial Semi-Bandits with Probabilistically Triggered Arms and Its Applications
<https://papers.nips.cc/paper/6716-improving-regret-bounds-for-combinatorial-semi-bandits-with-probabilistically-triggered-arms-and-its-applications>
AUTHORS: Qinshi Wang, Wei Chen
HIGHLIGHT: We resolve a serious issue in the prior CMAB-T studies where the regret bounds contain a possibly exponentially large factor of $1/p^*$, where p^* is the minimum positive probability that an arm is triggered by any action.
- 112, TITLE: Predictive-State Decoders: Encoding the Future into Recurrent Networks
<https://papers.nips.cc/paper/6717-predictive-state-decoders-encoding-the-future-into-recurrent-networks>
AUTHORS: Arun Venkatraman, Nicholas Rhinehart, Wen Sun, Lerrel Pinto, Martial Hebert, Byron Boots, Kris Kitani, J. Bagnell
HIGHLIGHT: We seek to combine the advantages of RNNs and PSRs by augmenting existing state-of-the-art recurrent neural networks with Predictive-State Decoders (PSDs), which add supervision to the network's internal state representation to target predicting future observations.
- 113, TITLE: Optimistic posterior sampling for reinforcement learning: worst-case regret bounds
<https://papers.nips.cc/paper/6718-optimistic-posterior-sampling-for-reinforcement-learning-worst-case-regret-bounds>
AUTHORS: Shipra Agrawal, Randy Jia
HIGHLIGHT: We present an algorithm based on posterior sampling (aka Thompson sampling) that achieves near-optimal worst-case regret bounds when the underlying Markov Decision Process (MDP) is communicating with a finite, though unknown, diameter.
- 114, TITLE: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results
<https://papers.nips.cc/paper/6719-mean-teachers-are-better-role-models-weight-averaged-consistency-targets-improve-semi-supervised-deep-learning-results>
AUTHORS: Antti Tarvainen, Harri Valpola
HIGHLIGHT: To overcome this problem, we propose Mean Teacher, a method that averages model weights instead of label predictions.
- 115, TITLE: Matching neural paths: transfer from recognition to correspondence search
<https://papers.nips.cc/paper/6720-matching-neural-paths-transfer-from-recognition-to-correspondence-search>
AUTHORS: Nikolay Savinov, Lubor Ladicky, Marc Pollefeys
HIGHLIGHT: In this work we focus on low-level correspondences --- a highly ambiguous matching problem.
- 116, TITLE: Linearly constrained Gaussian processes
<https://papers.nips.cc/paper/6721-linearly-constrained-gaussian-processes>
AUTHORS: Carl Jidling, Niklas Wahlstr?m, Adrian Wills, Thomas B. Sch?n
HIGHLIGHT: We consider a modification of the covariance function in Gaussian processes to correctly account for known linear constraints.
- 117, TITLE: Fixed-Rank Approximation of a Positive-Semidefinite Matrix from Streaming Data
<https://papers.nips.cc/paper/6722-fixed-rank-approximation-of-a-positive-semidefinite-matrix-from-streaming-data>
AUTHORS: Joel A. Tropp, Alp Yurtsever, Madeleine Udell, Volkan Cevher
HIGHLIGHT: This paper develops a new algorithm for fixed-rank psd approximation from a sketch.
- 118, TITLE: Multi-Modal Imitation Learning from Unstructured Demonstrations using Generative Adversarial Nets
<https://papers.nips.cc/paper/6723-multi-modal-imitation-learning-from-unstructured-demonstrations-using-generative-adversarial-nets>
AUTHORS: Karol Hausman, Yevgen Chebotar, Stefan Schaal, Gaurav Sukhatme, Joseph J. Lim
HIGHLIGHT: In this paper, we propose a multi-modal imitation learning framework that is able to segment and imitate skills from unlabelled and unstructured demonstrations by learning skill segmentation and imitation learning jointly.
- 119, TITLE: Learning to Inpaint for Image Compression

<https://papers.nips.cc/paper/6724-learning-to-inpaint-for-image-compression>

AUTHORS: Mohammad Haris Baig, Vladlen Koltun, Lorenzo Torresani

HIGHLIGHT: We present two architectural recipes in the context of multi-stage progressive encoders and empirically demonstrate their importance on compression performance.

120, TITLE: Adaptive Bayesian Sampling with Monte Carlo EM

<https://papers.nips.cc/paper/6725-adaptive-bayesian-sampling-with-monte-carlo-em>

AUTHORS: Anirban Roychowdhury, Srinivasan Parthasarathy

HIGHLIGHT: We present a novel technique for learning the mass matrices in samplers obtained from discretized dynamics that preserve some energy function.

121, TITLE: ADMM without a Fixed Penalty Parameter: Faster Convergence with New Adaptive Penalization

<https://papers.nips.cc/paper/6726-admm-without-a-fixed-penalty-parameter-faster-convergence-with-new-adaptive-penalization>

AUTHORS: Yi Xu, Mingrui Liu, Qihang Lin, Tianbao Yang

HIGHLIGHT: In this paper, we develop a new ADMM and its linearized variant with a new adaptive scheme to update the penalty parameter.

122, TITLE: Shape and Material from Sound

<https://papers.nips.cc/paper/6727-shape-and-material-from-sound>

AUTHORS: Zhoutong Zhang, Qiuja Li, Zhengjia Huang, Jiajun Wu, Josh Tenenbaum, Bill Freeman

HIGHLIGHT: In this paper, we build machines to approximate such competency.

123, TITLE: Flexible statistical inference for mechanistic models of neural dynamics

<https://papers.nips.cc/paper/6728-flexible-statistical-inference-for-mechanistic-models-of-neural-dynamics>

AUTHORS: Jan-Matthias Lueckmann, Pedro J. Goncalves, Giacomo Bassetto, Kaan ?cal, Marcel Nonnenmacher, Jakob H. Macke

HIGHLIGHT: We propose to overcome this limitation by using likelihood-free inference approaches (also known as Approximate Bayesian Computation, ABC) to perform full Bayesian inference on single-neuron models.

124, TITLE: Online Prediction with Selfish Experts

<https://papers.nips.cc/paper/6729-online-prediction-with-selfish-experts>

AUTHORS: Tim Roughgarden, Okke Schrijvers

HIGHLIGHT: We consider the problem of binary prediction with expert advice in settings where experts have agency and seek to maximize their credibility.

125, TITLE: Tensor Biclustering

<https://papers.nips.cc/paper/6730-tensor-biclustering>

AUTHORS: Soheil Feizi, Hamid Javadi, David Tse

HIGHLIGHT: We study the information-theoretic limit of this problem under a generative model.

126, TITLE: DPSCREEN: Dynamic Personalized Screening

<https://papers.nips.cc/paper/6731-dpscreen-dynamic-personalized-screening>

AUTHORS: Kartik Ahuja, William Zame, Mihaela van der Schaar

HIGHLIGHT: In this work, we develop the first framework to construct screening policies for a large class of disease models.

127, TITLE: Learning Unknown Markov Decision Processes: A Thompson Sampling Approach

<https://papers.nips.cc/paper/6732-learning-unknown-markov-decision-processes-a-thompson-sampling-approach>

AUTHORS: Yi Ouyang, Mukul Agrani, Ashutosh Nayyar, Rahul Jain

HIGHLIGHT: We propose a Thompson Sampling-based reinforcement learning algorithm with dynamic episodes (TSDE).

128, TITLE: Testing and Learning on Distributions with Symmetric Noise Invariance

<https://papers.nips.cc/paper/6733-testing-and-learning-on-distributions-with-symmetric-noise-invariance>

AUTHORS: Ho Chung Law, Christopher Yau, Dino Sejdinovic

HIGHLIGHT: We propose distances between distributions which encode invariance to additive symmetric noise, aimed at testing whether the assumed true underlying processes differ.

129, TITLE: A Dirichlet Mixture Model of Hawkes Processes for Event Sequence Clustering

<https://papers.nips.cc/paper/6734-a-dirichlet-mixture-model-of-hawkes-processes-for-event-sequence-clustering>

AUTHORS: Hongteng Xu, Hongyuan Zha

HIGHLIGHT: To solve this problem, we propose and discuss an effective model-based clustering method based on a novel Dirichlet mixture model of a special but significant type of point processes --- Hawkes process.

130, TITLE: Deanonymization in the Bitcoin P2P Network
<https://papers.nips.cc/paper/6735-deanonymization-in-the-bitcoin-p2p-network>
AUTHORS: Giulia Fanti, Pramod Viswanath
HIGHLIGHT: In this paper, we model the Bitcoin networking stack and analyze its anonymity properties, both pre- and post-2015.

131, TITLE: Accelerated consensus via Min-Sum Splitting
<https://papers.nips.cc/paper/6736-accelerated-consensus-via-min-sum-splitting>
AUTHORS: Patrick Rebeschini, Sekhar C. Tatikonda
HIGHLIGHT: We apply the Min-Sum message-passing protocol to solve the consensus problem in distributed optimization.

132, TITLE: Generalized Linear Model Regression under Distance-to-set Penalties
<https://papers.nips.cc/paper/6737-generalized-linear-model-regression-under-distance-to-set-penalties>
AUTHORS: Jason Xu, Eric Chi, Kenneth Lange
HIGHLIGHT: This paper explores instead penalizing the squared distance to constraint sets.

133, TITLE: Adaptive stimulus selection for optimizing neural population responses
<https://papers.nips.cc/paper/6738-adaptive-stimulus-selection-for-optimizing-neural-population-responses>
AUTHORS: Benjamin Cowley, Ryan Williamson, Katerina Clemens, Matthew Smith, Byron M. Yu
HIGHLIGHT: We propose "Adept," an adaptive stimulus selection method that can optimize population objective functions.

134, TITLE: Nonbacktracking Bounds on the Influence in Independent Cascade Models
<https://papers.nips.cc/paper/6739-nonbacktracking-bounds-on-the-influence-in-independent-cascade-models>
AUTHORS: Emmanuel Abbe, Sanjeev Kulkarni, Eun Jee Lee
HIGHLIGHT: This paper develops upper and lower bounds on the influence measure in a network, more precisely, the expected number of nodes that a seed set can influence in the independent cascade model.

135, TITLE: Learning with Feature Evolvable Streams
<https://papers.nips.cc/paper/6740-learning-with-feature-evolvable-streams>
AUTHORS: Bo-Jian Hou, Lijun Zhang, Zhi-Hua Zhou
HIGHLIGHT: In this paper, we propose a novel learning paradigm: Feature Evolvable Streaming Learning where old features would vanish and new features would occur.

136, TITLE: Online Convex Optimization with Stochastic Constraints
<https://papers.nips.cc/paper/6741-online-convex-optimization-with-stochastic-constraints>
AUTHORS: Hao Yu, Michael Neely, Xiaohan Wei
HIGHLIGHT: To solve this problem, this paper proposes a new algorithm that achieves $\mathcal{O}(\sqrt{T})$ expected regret and constraint violations and $\mathcal{O}(\sqrt{T} \log(T))$ high probability regret and constraint violations.

137, TITLE: Max-Margin Invariant Features from Transformed Unlabelled Data
<https://papers.nips.cc/paper/6742-max-margin-invariant-features-from-transformed-unlabelled-data>
AUTHORS: Dipan Pal, Ashwin Kannan, Gautam Arakalgud, Marios Savvides
HIGHLIGHT: In this paper, we study kernels that are invariant to a unitary group while having theoretical guarantees in addressing the important practical issue of unavailability of transformed versions of labelled data.

138, TITLE: Regularized Modal Regression with Applications in Cognitive Impairment Prediction
<https://papers.nips.cc/paper/6743-regularized-modal-regression-with-applications-in-cognitive-impairment-prediction>
AUTHORS: Xiaoqian Wang, Hong Chen, Weidong Cai, Dinggang Shen, Heng Huang
HIGHLIGHT: In this paper, we go beyond this criterion by investigating the regularized modal regression from a statistical learning viewpoint.

139, TITLE: Translation Synchronization via Truncated Least Squares
<https://papers.nips.cc/paper/6744-translation-synchronization-via-truncated-least-squares>
AUTHORS: Xiangru Huang, Zhenxiao Liang, Chandrajit Bajaj, Qixing Huang
HIGHLIGHT: In this paper, we introduce a robust algorithm, `\textsl{TranSync}`, for the 1D translation synchronization problem, in which the aim is to recover the global coordinates of a set of nodes from noisy measurements of relative coordinates along an observation graph.

140, TITLE: From which world is your graph
<https://papers.nips.cc/paper/6745-from-which-world-is-your-graph>
AUTHORS: Cheng Li, Felix MF Wong, Zhenming Liu, Varun Kanade
HIGHLIGHT: This work focuses on unifying two of the most widely used link-formation models: the stochastic block model (SBM) and the small world (or latent space) model (SWM).

141, TITLE: A New Alternating Direction Method for Linear Programming
<https://papers.nips.cc/paper/6746-a-new-alternating-direction-method-for-linear-programming>
AUTHORS: Sinong Wang, Ness Shroff
HIGHLIGHT: In this paper, we propose a new variable splitting method of LP and prove that our method has a convergence rate of $O(\|\mathbf{A}\|^2 \log(1/\epsilon))$.

142, TITLE: Regret Analysis for Continuous Dueling Bandit
<https://papers.nips.cc/paper/6747-regret-analysis-for-continuous-dueling-bandit>
AUTHORS: Wataru Kumagai
HIGHLIGHT: In this paper, we address a dueling bandit problem based on a cost function over a continuous space.

143, TITLE: Best Response Regression
<https://papers.nips.cc/paper/6748-best-response-regression>
AUTHORS: Omer Ben-Porat, Moshe Tennenholtz
HIGHLIGHT: In this work, we initiate the study of strategic predictions in machine learning.

144, TITLE: TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning
<https://papers.nips.cc/paper/6749-terngrad-ternary-gradients-to-reduce-communication-in-distributed-deep-learning>
AUTHORS: Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, Hai Li
HIGHLIGHT: In this work, we propose TernGrad that uses ternary gradients to accelerate distributed deep learning in data parallelism.

145, TITLE: Learning Affinity via Spatial Propagation Networks
<https://papers.nips.cc/paper/6750-learning-affinity-via-spatial-propagation-networks>
AUTHORS: Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, Jan Kautz
HIGHLIGHT: In this paper, we propose a spatial propagation networks for learning affinity matrix.

146, TITLE: Linear regression without correspondence
<https://papers.nips.cc/paper/6751-linear-regression-without-correspondence>
AUTHORS: Daniel J. Hsu, Kevin Shi, Xiaorui Sun
HIGHLIGHT: This article considers algorithmic and statistical aspects of linear regression when the correspondence between the covariates and the responses is unknown.

147, TITLE: NeuralFDR: Learning Discovery Thresholds from Hypothesis Features
<https://papers.nips.cc/paper/6752-neuralfdr-learning-discovery-thresholds-from-hypothesis-features>
AUTHORS: Fei Xia, Martin J. Zhang, James Y. Zou, David Tse
HIGHLIGHT: We propose a new algorithm, NeuralFDR, which automatically learns a discovery threshold as a function of all the hypothesis features.

148, TITLE: Cost efficient gradient boosting
<https://papers.nips.cc/paper/6753-cost-efficient-gradient-boosting>
AUTHORS: Sven Peter, Ferran Diego, Fred A. Hamprecht, Boaz Nadler
HIGHLIGHT: In this work we propose a budget-aware strategy based on deep boosted regression trees.

149, TITLE: Probabilistic Rule Realization and Selection
<https://papers.nips.cc/paper/6754-probabilistic-rule-realization-and-selection>
AUTHORS: Haizi Yu, Tianxi Li, Lav R. Varshney
HIGHLIGHT: Under a probabilistic setting for discrete input spaces, we focus on the rule realization problem which generates input sample distributions that follow the given rules.

150, TITLE: Nearest-Neighbor Sample Compression: Efficiency, Consistency, Infinite Dimensions
<https://papers.nips.cc/paper/6755-nearest-neighbor-sample-compression-efficiency-consistency-infinite-dimensions>

- AUTHORS: Aryeh Kontorovich, Sivan Sabato, Roi Weiss
HIGHLIGHT: We pose several challenging open problems for future research.
- 151, TITLE: A Scale Free Algorithm for Stochastic Bandits with Bounded Kurtosis
<https://papers.nips.cc/paper/6756-a-scale-free-algorithm-for-stochastic-bandits-with-bounded-kurtosis>
AUTHORS: Tor Lattimore
HIGHLIGHT: A Scale Free Algorithm for Stochastic Bandits with Bounded Kurtosis
- 152, TITLE: Learning Multiple Tasks with Multilinear Relationship Networks
<https://papers.nips.cc/paper/6757-learning-multiple-tasks-with-multilinear-relationship-networks>
AUTHORS: Mingsheng Long, ZHANGJIE CAO, Jianmin Wang, Philip S. Yu
HIGHLIGHT: This paper presents Multilinear Relationship Networks (MRN) that discover the task relationships based on novel tensor normal priors over parameter tensors of multiple task-specific layers in deep convolutional networks.
- 153, TITLE: Deep Hyperalignment
<https://papers.nips.cc/paper/6758-deep-hyperalignment>
AUTHORS: Muhammad Yousefnezhad, Daoqiang Zhang
HIGHLIGHT: This paper proposes Deep Hyperalignment (DHA) as a regularized, deep extension, scalable Hyperalignment (HA) method, which is well-suited for applying functional alignment to fMRI datasets with nonlinearity, high-dimensionality (broad ROI), and a large number of subjects.
- 154, TITLE: Online to Offline Conversions, Universality and Adaptive Minibatch Sizes
<https://papers.nips.cc/paper/6759-online-to-offline-conversions-universality-and-adaptive-minibatch-sizes>
AUTHORS: Kfir Levy
HIGHLIGHT: We present an approach towards convex optimization that relies on a novel scheme which converts adaptive online algorithms into offline methods.
- 155, TITLE: Stochastic Optimization with Variance Reduction for Infinite Datasets with Finite Sum Structure
<https://papers.nips.cc/paper/6760-stochastic-optimization-with-variance-reduction-for-infinite-datasets-with-finite-sum-structure>
AUTHORS: Alberto Bietti, Julien Mairal
HIGHLIGHT: In this paper, we introduce a variance reduction approach for these settings when the objective is composite and strongly convex.
- 156, TITLE: Deep Learning with Topological Signatures
<https://papers.nips.cc/paper/6761-deep-learning-with-topological-signatures>
AUTHORS: Christoph Hofer, Roland Kwitt, Marc Niethammer, Andreas Uhl
HIGHLIGHT: In contrast, we propose a technique that enables us to input topological signatures to deep neural networks and learn a task-optimal representation during training.
- 157, TITLE: Predicting User Activity Level In Point Processes With Mass Transport Equation
<https://papers.nips.cc/paper/6762-predicting-user-activity-level-in-point-processes-with-mass-transport-equation>
AUTHORS: Yichen Wang, Xiaojing Ye, Hongyuan Zha, Le Song
HIGHLIGHT: In this paper, we propose a framework that provides an unbiased estimator of the probability mass function of point processes.
- 158, TITLE: Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues
<https://papers.nips.cc/paper/6763-submultiplicative-glivenko-cantelli-and-uniform-convergence-of-revenues>
AUTHORS: Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, Shay Moran, Amir Yehudayoff
HIGHLIGHT: In this work we derive a variant of the classic Glivenko-Cantelli Theorem, which asserts uniform convergence of the empirical Cumulative Distribution Function (CDF) to the CDF of the underlying distribution.
- 159, TITLE: Deep Dynamic Poisson Factorization Model
<https://papers.nips.cc/paper/6764-deep-dynamic-poisson-factorization-model>
AUTHORS: Chengyue Gong, win-bin huang
HIGHLIGHT: A new model, named as deep dynamic poisson factorization model, is proposed in this paper for analyzing sequential count vectors.
- 160, TITLE: Positive-Unlabeled Learning with Non-Negative Risk Estimator
<https://papers.nips.cc/paper/6765-positive-unlabeled-learning-with-non-negative-risk-estimator>
AUTHORS: Ryuichi Kiryo, Gang Niu, Marthinus C. du Plessis, Masashi Sugiyama

HIGHLIGHT: In this paper, we propose a non-negative risk estimator for PU learning: when getting minimized, it is more robust against overfitting, and thus we are able to use very flexible models (such as deep neural networks) given limited P data.

161, **TITLE:** Optimal Sample Complexity of M-wise Data for Top-K Ranking
<https://papers.nips.cc/paper/6766-optimal-sample-complexity-of-m-wise-data-for-top-k-ranking>
AUTHORS: Minje Jang, Sunghyun Kim, Changho Suh, Sewoong Oh
HIGHLIGHT: We explore the top-K rank aggregation problem in which one aims to recover a consistent ordering that focuses on top-K ranked items based on partially revealed preference information.

162, **TITLE:** Reliable Decision Support using Counterfactual Models
<https://papers.nips.cc/paper/6767-reliable-decision-support-using-counterfactual-models>
AUTHORS: Peter Schulam, Suchi Saria
HIGHLIGHT: We propose using a different learning objective that predicts counterfactuals instead of predicting outcomes under an existing action policy as in supervised learning.

163, **TITLE:** QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding
<https://papers.nips.cc/paper/6768-qsgd-communication-efficient-sgd-via-gradient-quantization-and-encoding>
AUTHORS: Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, Milan Vojnovic
HIGHLIGHT: In this paper, we propose Quantized SGD (QSGD), a family of compression schemes for gradient updates which provides convergence guarantees.

164, **TITLE:** Convergent Block Coordinate Descent for Training Tikhonov Regularized Deep Neural Networks
<https://papers.nips.cc/paper/6769-convergent-block-coordinate-descent-for-training-tikhonov-regularized-deep-neural-networks>
AUTHORS: Ziming Zhang, Matthew Brand
HIGHLIGHT: By lifting the ReLU function into a higher dimensional space, we develop a smooth multi-convex formulation for training feed-forward deep neural networks (DNNs).

165, **TITLE:** Train longer, generalize better: closing the generalization gap in large batch training of neural networks
<https://papers.nips.cc/paper/6770-train-longer-generalize-better-closing-the-generalization-gap-in-large-batch-training-of-neural-networks>
AUTHORS: Elad Hoffer, Itay Hubara, Daniel Soudry
HIGHLIGHT: We therefore propose a "random walk on a random landscape" statistical model which is known to exhibit similar "ultra-slow" diffusion behavior.

166, **TITLE:** Flexpoint: An Adaptive Numerical Format for Efficient Training of Deep Neural Networks
<https://papers.nips.cc/paper/6771-flexpoint-an-adaptive-numerical-format-for-efficient-training-of-deep-neural-networks>
AUTHORS: Urs Koster, Tristan Webb, Xin Wang, Marcel Nassar, Arjun K. Bansal, William Constable, Oguz Elibol, Scott Gray, Stewart Hall, Luke Hornof, Amir Khosrowshahi, Carey Kloss, Ruby J. Pai, Naveen Rao
HIGHLIGHT: Here we present the Flexpoint data format, aiming at a complete replacement of 32-bit floating point format training and inference, designed to support modern deep network topologies without modifications.

167, **TITLE:** Model evidence from nonequilibrium simulations
<https://papers.nips.cc/paper/6772-model-evidence-from-nonequilibrium-simulations>
AUTHORS: Michael Habeck
HIGHLIGHT: We introduce estimators for the model evidence that combine forward and backward simulations and show for various challenging models that the evidence estimators outperform forward and reverse AIS.

168, **TITLE:** Minimal Exploration in Structured Stochastic Bandits
<https://papers.nips.cc/paper/6773-minimal-exploration-in-structured-stochastic-bandits>
AUTHORS: Richard Combes, Stefan Magureanu, Alexandre Proutiere
HIGHLIGHT: This paper introduces and addresses a wide class of stochastic bandit problems where the function mapping the arm to the corresponding reward exhibits some known structural properties.

169, **TITLE:** Learned D-AMP: Principled Neural Network based Compressive Image Recovery
<https://papers.nips.cc/paper/6774-learned-d-amp-principled-neural-network-based-compressive-image-recovery>
AUTHORS: Chris Metzler, Ali Mousavi, Richard Baraniuk
HIGHLIGHT: Taking inspiration from this work, we develop a novel neural network architecture that mimics the behavior of the denoising-based approximate message passing (D-AMP) algorithm.

170, **TITLE:** Deliberation Networks: Sequence Generation Beyond One-Pass Decoding

- <https://papers.nips.cc/paper/6775-deliberation-networks-sequence-generation-beyond-one-pass-decoding>
AUTHORS: Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, Tie-Yan Liu
HIGHLIGHT: In this work, we introduce the deliberation process into the encoder-decoder framework and propose deliberation networks for sequence generation.
- 171, TITLE: Adaptive Clustering through Semidefinite Programming
<https://papers.nips.cc/paper/6776-adaptive-clustering-through-semidefinite-programming>
AUTHORS: Martin Royer
HIGHLIGHT: We analyze the clustering problem through a flexible probabilistic model that aims to identify an optimal partition on the sample X_1, \dots, X_n .
- 172, TITLE: Log-normality and Skewness of Estimated State/Action Values in Reinforcement Learning
<https://papers.nips.cc/paper/6777-log-normality-and-skewness-of-estimated-stateaction-values-in-reinforcement-learning>
AUTHORS: Liangpeng Zhang, Ke Tang, Xin Yao
HIGHLIGHT: In this paper, we show that a state/action value estimated using the Bellman equation can be decomposed to a weighted sum of path-wise values that follow log-normal distributions.
- 173, TITLE: Repeated Inverse Reinforcement Learning
<https://papers.nips.cc/paper/6778-repeated-inverse-reinforcement-learning>
AUTHORS: Kareem Amin, Nan Jiang, Satinder Singh
HIGHLIGHT: We introduce a novel repeated Inverse Reinforcement Learning problem: the agent has to act on behalf of a human in a sequence of tasks and wishes to minimize the number of tasks that it surprises the human by acting suboptimally with respect to how the human would have acted.
- 174, TITLE: The Numerics of GANs
<https://papers.nips.cc/paper/6779-the-numeric-of-gans>
AUTHORS: Lars Mescheder, Sebastian Nowozin, Andreas Geiger
HIGHLIGHT: In this paper, we analyze the numerics of common algorithms for training Generative Adversarial Networks (GANs).
- 175, TITLE: Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search
<https://papers.nips.cc/paper/6780-practical-bayesian-optimization-for-model-fitting-with-bayesian-adaptive-direct-search>
AUTHORS: Luigi Acerbi, Wei Ji
HIGHLIGHT: Computational models in fields such as computational neuroscience are often evaluated via stochastic simulation or numerical approximation.
- 176, TITLE: Learning Chordal Markov Networks via Branch and Bound
<https://papers.nips.cc/paper/6781-learning-chordal-markov-networks-via-branch-and-bound>
AUTHORS: Kari Rantanen, Antti Hyttinen, Matti J?rvisalo
HIGHLIGHT: We present a new algorithmic approach for the task of finding a chordal Markov network structure that maximizes a given scoring function.
- 177, TITLE: Revenue Optimization with Approximate Bid Predictions
<https://papers.nips.cc/paper/6782-revenue-optimization-with-approximate-bid-predictions>
AUTHORS: Andres Munoz, Sergei Vassilvitskii
HIGHLIGHT: In this work, we show how to reduce reserve price optimization to the standard setting of prediction under squared loss, a well understood problem in the learning community.
- 178, TITLE: Solving Most Systems of Random Quadratic Equations
<https://papers.nips.cc/paper/6783-solving-most-systems-of-random-quadratic-equations>
AUTHORS: Gang Wang, Georgios Giannakis, Yousef Saad, Jie Chen
HIGHLIGHT: This paper deals with finding an n -dimensional solution \mathbf{x} to a system of quadratic equations $\|y_i - \langle \mathbf{a}_i, \mathbf{x} \rangle\|^2, \forall i \in [m]$, which in general is known to be NP-hard.
- 179, TITLE: Unsupervised Learning of Disentangled and Interpretable Representations from Sequential Data
<https://papers.nips.cc/paper/6784-unsupervised-learning-of-disentangled-and-interpretable-representations-from-sequential-data>
AUTHORS: Wei-Ning Hsu, Yu Zhang, James Glass
HIGHLIGHT: We present a factorized hierarchical variational autoencoder, which learns disentangled and interpretable representations from sequential data without supervision.

- 180, TITLE: Lookahead Bayesian Optimization with Inequality Constraints
<https://papers.nips.cc/paper/6785-lookahead-bayesian-optimization-with-inequality-constraints>
AUTHORS: Remi Lam, Karen Willcox
HIGHLIGHT: To address this limitation, we propose a lookahead approach that selects the next evaluation in order to maximize the long-term feasible reduction of the objective function.
- 181, TITLE: Hierarchical Methods of Moments
<https://papers.nips.cc/paper/6786-hierarchical-methods-of-moments>
AUTHORS: Matteo Ruffini, Guillaume Rabusseau, Borja Balle
HIGHLIGHT: In this paper we present a hierarchical approach to methods of moments to circumvent such limitations.
- 182, TITLE: Interpretable and Globally Optimal Prediction for Textual Grounding using Image Concepts
<https://papers.nips.cc/paper/6787-interpretable-and-globally-optimal-prediction-for-textual-grounding-using-image-concepts>
AUTHORS: Raymond Yeh, Jinjun Xiong, Wen-Mei Hwu, Minh Do, Alexander Schwing
HIGHLIGHT: In this work, we demonstrate that we can cast the problem of textual grounding into a unified framework that permits efficient search over all possible bounding boxes.
- 183, TITLE: Revisit Fuzzy Neural Network: Demystifying Batch Normalization and ReLU with Generalized Hamming Network
<https://papers.nips.cc/paper/6788-revisit-fuzzy-neural-network-demystifying-batch-normalization-and-relu-with-generalized-hamming-network>
AUTHORS: Lixin Fan
HIGHLIGHT: We revisit fuzzy neural network with a cornerstone notion of generalized hamming distance, which provides a novel and theoretically justified framework to re-interpret many useful neural network techniques in terms of fuzzy logic.
- 184, TITLE: Speeding Up Latent Variable Gaussian Graphical Model Estimation via Nonconvex Optimization
<https://papers.nips.cc/paper/6789-speeding-up-latent-variable-gaussian-graphical-model-estimation-via-nonconvex-optimization>
AUTHORS: Pan Xu, Jian Ma, Quanquan Gu
HIGHLIGHT: In order to speed up the estimation of the sparse plus low-rank components, we propose a sparsity constrained maximum likelihood estimator based on matrix factorization and an efficient alternating gradient descent algorithm with hard thresholding to solve it.
- 185, TITLE: Batch Renormalization: Towards Reducing Minibatch Dependence in Batch-Normalized Models
<https://papers.nips.cc/paper/6790-batch-renormalization-towards-reducing-minibatch-dependence-in-batch-normalized-models>
AUTHORS: Sergey Ioffe
HIGHLIGHT: We propose Batch Renormalization, a simple and effective extension to ensure that the training and inference models generate the same outputs that depend on individual examples rather than the entire minibatch.
- 186, TITLE: Generating steganographic images via adversarial training
<https://papers.nips.cc/paper/6791-generating-steganographic-images-via-adversarial-training>
AUTHORS: Jamie Hayes, George Danezis
HIGHLIGHT: In this paper, we apply adversarial training techniques to the discriminative task of learning a steganographic algorithm.
- 187, TITLE: Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration
<https://papers.nips.cc/paper/6792-near-linear-time-approximation-algorithms-for-optimal-transport-via-sinkhorn-iteration>
AUTHORS: Jason Altschuler, Jonathan Weed, Philippe Rigollet
HIGHLIGHT: This paper demonstrates that this ambitious goal is in fact achieved by Cuturi's Sinkhorn Distances.
- 188, TITLE: PixelGAN Autoencoders
<https://papers.nips.cc/paper/6793-pixelgan-autoencoders>
AUTHORS: Alireza Makhzani, Brendan J. Frey
HIGHLIGHT: In this paper, we describe the "PixelGAN autoencoder", a generative autoencoder in which the generative path is a convolutional autoregressive neural network on pixels (PixelCNN) that is conditioned on a latent code, and the recognition path uses a generative adversarial network (GAN) to impose a prior distribution on the latent code.
- 189, TITLE: Consistent Multitask Learning with Nonlinear Output Relations
<https://papers.nips.cc/paper/6794-consistent-multitask-learning-with-nonlinear-output-relations>
AUTHORS: Carlo Ciliberto, Alessandro Rudi, Lorenzo Rosasco, Massimiliano Pontil
HIGHLIGHT: In this paper, we tackle this issue by casting the problem within the framework of structured prediction.

- 190, TITLE: Alternating minimization for dictionary learning with random initialization
<https://papers.nips.cc/paper/6795-alternating-minimization-for-dictionary-learning-with-random-initialization>
AUTHORS: Niladri Chatterji, Peter L. Bartlett
HIGHLIGHT: We present theoretical guarantees for an alternating minimization algorithm for the dictionary learning/sparse coding problem.
- 191, TITLE: Learning ReLUs via Gradient Descent
<https://papers.nips.cc/paper/6796-learning-relus-via-gradient-descent>
AUTHORS: Mahdi Soltanolkotabi
HIGHLIGHT: In this paper we study the problem of learning Rectified Linear Units (ReLUs) which are functions of the form $\|\text{vct}\{x\}\mapsto \max(0, \lvert \text{vct}\{w\}, \text{vct}\{x\} \rvert)$ with $\text{vct}\{w\} \in \mathbb{R}^d$ denoting the weight vector.
- 192, TITLE: Stabilizing Training of Generative Adversarial Networks through Regularization
<https://papers.nips.cc/paper/6797-stabilizing-training-of-generative-adversarial-networks-through-regularization>
AUTHORS: Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, Thomas Hofmann
HIGHLIGHT: We overcome this fundamental limitation and propose a new regularization approach with low computational cost that yields a stable GAN training procedure.
- 193, TITLE: Expectation Propagation with Stochastic Kinetic Model in Complex Interaction Systems
<https://papers.nips.cc/paper/6798-expectation-propagation-with-stochastic-kinetic-model-in-complex-interaction-systems>
AUTHORS: Le Fang, Fan Yang, Wen Dong, Tong Guan, Chunming Qiao
HIGHLIGHT: In this paper, we formulate the dynamics of a complex interacting network as a stochastic process driven by a sequence of events, and develop expectation propagation algorithms to make inferences from noisy observations.
- 194, TITLE: Data-Efficient Reinforcement Learning in Continuous State-Action Gaussian-POMDPs
<https://papers.nips.cc/paper/6799-data-efficient-reinforcement-learning-in-continuous-state-action-gaussian-pomdps>
AUTHORS: Rowan McAllister, Carl Edward Rasmussen
HIGHLIGHT: We present a data-efficient reinforcement learning method for continuous state-action systems under significant observation noise.
- 195, TITLE: Compatible Reward Inverse Reinforcement Learning
<https://papers.nips.cc/paper/6800-compatible-reward-inverse-reinforcement-learning>
AUTHORS: Alberto Maria Metelli, Matteo Pirota, Marcello Restelli
HIGHLIGHT: After introducing our approach for finite domains, we extend it to continuous ones.
- 196, TITLE: First-Order Adaptive Sample Size Methods to Reduce Complexity of Empirical Risk Minimization
<https://papers.nips.cc/paper/6801-first-order-adaptive-sample-size-methods-to-reduce-complexity-of-empirical-risk-minimization>
AUTHORS: Aryan Mokhtari, Alejandro Ribeiro
HIGHLIGHT: The sample size is then grown geometrically -- e.g., scaling by a factor of two -- and use the solution of the previous ERM as a warm start for the new ERM.
- 197, TITLE: Hiding Images in Plain Sight: Deep Steganography
<https://papers.nips.cc/paper/6802-hiding-images-in-plain-sight-deep-steganography>
AUTHORS: Shumeet Baluja
HIGHLIGHT: In this study, we attempt to place a full size color image within another image of the same size.
- 198, TITLE: Neural Program Meta-Induction
<https://papers.nips.cc/paper/6803-neural-program-meta-induction>
AUTHORS: Jacob Devlin, Rudy R. Bunel, Rishabh Singh, Matthew Hausknecht, Pushmeet Kohli
HIGHLIGHT: Specifically, we propose two novel approaches for cross-task knowledge transfer to improve program induction in limited-data scenarios.
To test the efficacy of our methods, we constructed a new benchmark of programs written in the Karel programming language.
- 199, TITLE: Bayesian Dyadic Trees and Histograms for Regression
<https://papers.nips.cc/paper/6804-bayesian-dyadic-trees-and-histograms-for-regression>
AUTHORS: Stephanie van der Pas, Veronika Rockova
HIGHLIGHT: In this work, we shed light on the machinery behind Bayesian variants of these methods.
- 200, TITLE: A graph-theoretic approach to multitasking

<https://papers.nips.cc/paper/6805-a-graph-theoretic-approach-to-multitasking>
AUTHORS: Noga Alon, Daniel Reichman, Igor Shinkar, Tal Wagner, Sebastian Musslick, Jonathan D. Cohen, Tom Griffiths, Biswadip dey, Kayhan Ozcimder
HIGHLIGHT: In this paper we use a graph-theoretic analysis of network architecture to address this question, where tasks are represented as edges in a bipartite graph $G=(A \cup B, E)$.

201, TITLE: Consistent Robust Regression
<https://papers.nips.cc/paper/6806-consistent-robust-regression>
AUTHORS: Kush Bhatia, Prateek Jain, Parameswaran Kamalaruban, Purushottam Kar
HIGHLIGHT: In this work we present such an estimator, called CRR.

202, TITLE: Natural Value Approximators: Learning when to Trust Past Estimates
<https://papers.nips.cc/paper/6807-natural-value-approximators-learning-when-to-trust-past-estimates>
AUTHORS: Zhongwen Xu, Joseph Modayil, Hado P. van Hasselt, Andre Barreto, David Silver, Tom Schaul
HIGHLIGHT: We propose a mechanism that learns an interpolation between a direct value estimate and a projected value estimate computed from the encountered reward and the previous estimate.

203, TITLE: Bandits Dueling on Partially Ordered Sets
<https://papers.nips.cc/paper/6808-bandits-dueling-on-partially-ordered-sets>
AUTHORS: Julien Audiffren, Liva Ralaivola
HIGHLIGHT: We propose an algorithm, UnchainedBandits, that efficiently finds the set of optimal arms, or Pareto front, of any poset even when pairs of comparable arms cannot be a priori distinguished from pairs of incomparable arms, with a set of minimal assumptions.

204, TITLE: Elementary Symmetric Polynomials for Optimal Experimental Design
<https://papers.nips.cc/paper/6809-elementary-symmetric-polynomials-for-optimal-experimental-design>
AUTHORS: Zelda E. Mariet, Suvrit Sra
HIGHLIGHT: Specifically, we introduce models based on elementary symmetric polynomials; these polynomials capture "partial volumes" and offer a graded interpolation between the widely used A-optimal and D-optimal design models, obtaining each of them as special cases.

205, TITLE: Emergence of Language with Multi-agent Games: Learning to Communicate with Sequences of Symbols
<https://papers.nips.cc/paper/6810-emergence-of-language-with-multi-agent-games-learning-to-communicate-with-sequences-of-symbols>
AUTHORS: Serhii Havrylov, Ivan Titov
HIGHLIGHT: We compare a reinforcement learning approach and one using a differentiable relaxation (straight-through Gumbel-softmax estimator) and observe that the latter is much faster to converge and it results in more effective protocols.

206, TITLE: Training Deep Networks without Learning Rates Through Coin Betting
<https://papers.nips.cc/paper/6811-training-deep-networks-without-learning-rates-through-coin-betting>
AUTHORS: Francesco Orabona, Tatiana Tommasi
HIGHLIGHT: Instead, we reduce the optimization process to a game of betting on a coin and propose a learning rate free optimal algorithm for this scenario.

207, TITLE: Pixels to Graphs by Associative Embedding
<https://papers.nips.cc/paper/6812-pixels-to-graphs-by-associative-embedding>
AUTHORS: Alejandro Newell, Jia Deng
HIGHLIGHT: We present a method for training a convolutional neural network such that it takes in an input image and produces a full graph definition.

208, TITLE: Runtime Neural Pruning
<https://papers.nips.cc/paper/6813-runtime-neural-pruning>
AUTHORS: Ji Lin, Yongming Rao, Jiwen Lu, Jie Zhou
HIGHLIGHT: In this paper, we propose a Runtime Neural Pruning (RNP) framework which prunes the deep neural network dynamically at the runtime.

209, TITLE: Eigenvalue Decay Implies Polynomial-Time Learnability for Neural Networks
<https://papers.nips.cc/paper/6814-eigenvalue-decay-implies-polynomial-time-learnability-for-neural-networks>
AUTHORS: Surbhi Goel, Adam Klivans

HIGHLIGHT: In this work we show that a natural distributional assumption corresponding to $\{\text{eigenvalue decay}\}$ of the Gram matrix yields polynomial-time algorithms in the non-realizable setting for expressive classes of networks (e.g. feed-forward networks of ReLUs).

210, **TITLE:** MMD GAN: Towards Deeper Understanding of Moment Matching Network
<https://papers.nips.cc/paper/6815-mmd-gan-towards-deeper-understanding-of-moment-matching-network>
AUTHORS: Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, Barnabas Poczos
HIGHLIGHT: In this paper, we propose to improve both the model expressiveness of GMMN and its computational efficiency by introducing $\{\text{adversarial kernel learning}\}$ techniques, as the replacement of a fixed Gaussian kernel in the original GMMN.

211, **TITLE:** The Reversible Residual Network: Backpropagation Without Storing Activations
<https://papers.nips.cc/paper/6816-the-reversible-residual-network-backpropagation-without-storing-activations>
AUTHORS: Aidan N. Gomez, Mengye Ren, Raquel Urtasun, Roger B. Grosse
HIGHLIGHT: In this work, we present the Reversible Residual Network (RevNet), a variant of ResNets where each layer's activations can be reconstructed exactly from the next layer's.

212, **TITLE:** Fast Rates for Bandit Optimization with Upper-Confidence Frank-Wolfe
<https://papers.nips.cc/paper/6817-fast-rates-for-bandit-optimization-with-upper-confidence-frank-wolfe>
AUTHORS: Quentin Berthet, Vianney Perchet
HIGHLIGHT: We consider the problem of bandit optimization, inspired by stochastic optimization and online learning problems with bandit feedback.

213, **TITLE:** Zap Q-Learning
<https://papers.nips.cc/paper/6818-zap-q-learning>
AUTHORS: Adithya M Devraj, Sean Meyn
HIGHLIGHT: The Zap Q-learning algorithm introduced in this paper is an improvement of Watkins' original algorithm and recent competitors in several respects.

214, **TITLE:** Expectation Propagation for t-Exponential Family Using q-Algebra
<https://papers.nips.cc/paper/6819-expectation-propagation-for-t-exponential-family-using-q-algebra>
AUTHORS: Futoshi Futami, Issei Sato, Masashi Sugiyama
HIGHLIGHT: In this paper, we borrow the mathematical tools of q-algebra from statistical physics and show that the pseudo additivity of distributions allows us to perform calculation of t-exponential family distributions through natural parameters.

215, **TITLE:** Few-Shot Learning Through an Information Retrieval Lens
<https://papers.nips.cc/paper/6820-few-shot-learning-through-an-information-retrieval-lens>
AUTHORS: Eleni Triantafillou, Richard Zemel, Raquel Urtasun
HIGHLIGHT: We propose an information retrieval-inspired approach for this problem that is motivated by the increased importance of maximally leveraging all the available information in this low-data regime.

216, **TITLE:** Formal Guarantees on the Robustness of a Classifier against Adversarial Manipulation
<https://papers.nips.cc/paper/6821-formal-guarantees-on-the-robustness-of-a-classifier-against-adversarial-manipulation>
AUTHORS: Matthias Hein, Maksym Andriushchenko
HIGHLIGHT: We show in this paper for the first time formal guarantees on the robustness of a classifier by giving instance-specific $\{\text{lower bounds}\}$ on the norm of the input manipulation required to change the classifier decision.

217, **TITLE:** Associative Embedding: End-to-End Learning for Joint Detection and Grouping
<https://papers.nips.cc/paper/6822-associative-embedding-end-to-end-learning-for-joint-detection-and-grouping>
AUTHORS: Alejandro Newell, Zhiao Huang, Jia Deng
HIGHLIGHT: We introduce associative embedding, a novel method for supervising convolutional neural networks for the task of detection and grouping.

218, **TITLE:** Practical Locally Private Heavy Hitters
<https://papers.nips.cc/paper/6823-practical-locally-private-heavy-hitters>
AUTHORS: Raef Bassily, Kobbi Nissim, Uri Stemmer, Abhradeep Guha Thakurta
HIGHLIGHT: We present new practical local differentially private heavy hitters algorithms achieving optimal or near-optimal worst-case error -- TreeHist and Bitstogram.

219, **TITLE:** Large-Scale Quadratically Constrained Quadratic Program via Low-Discrepancy Sequences
<https://papers.nips.cc/paper/6824-large-scale-quadratically-constrained-quadratic-program-via-low-discrepancy-sequences>

AUTHORS: Kinjal Basu, Ankan Saha, Shaunak Chatterjee
HIGHLIGHT: In this paper, we develop a method that transforms the quadratic constraint into a linear form by a sampling a set of low-discrepancy points.

220, TITLE: Inhomogeneous Hypergraph Clustering with Applications
<https://papers.nips.cc/paper/6825-inhomogeneous-hypergraph-clustering-with-applications>
AUTHORS: Pan Li, Olgica Milenkovic
HIGHLIGHT: We hence propose a new hypergraph clustering technique, termed inhomogeneous hypergraph partitioning, which assigns different costs to different hyperedge cuts.

221, TITLE: Differentiable Learning of Logical Rules for Knowledge Base Reasoning
<https://papers.nips.cc/paper/6826-differentiable-learning-of-logical-rules-for-knowledge-base-reasoning>
AUTHORS: Fan Yang, Zhilin Yang, William W. Cohen
HIGHLIGHT: We propose a framework, Neural Logic Programming, that combines the parameter and structure learning of first-order logical rules in an end-to-end differentiable model.

222, TITLE: Deep Multi-task Gaussian Processes for Survival Analysis with Competing Risks
<https://papers.nips.cc/paper/6827-deep-multi-task-gaussian-processes-for-survival-analysis-with-competing-risks>
AUTHORS:
HIGHLIGHT: We propose a variational inference algorithm that is capable of learning the model parameters from time-to-event data while handling right censoring.

223, TITLE: Masked Autoregressive Flow for Density Estimation
<https://papers.nips.cc/paper/6828-masked-autoregressive-flow-for-density-estimation>
AUTHORS: George Papamakarios, Theo Pavlakou, Iain Murray
HIGHLIGHT: We describe an approach for increasing the flexibility of an autoregressive model, based on modelling the random numbers that the model uses internally when generating data.

224, TITLE: Non-convex Finite-Sum Optimization Via SCSG Methods
<https://papers.nips.cc/paper/6829-non-convex-finite-sum-optimization-via-scsg-methods>
AUTHORS: Lihua Lei, Cheng Ju, Jianbo Chen, Michael I. Jordan
HIGHLIGHT: We develop a class of algorithms, as variants of the stochastically controlled stochastic gradient (SCSG) methods, for the smooth nonconvex finite-sum optimization problem.

225, TITLE: Beyond normality: Learning sparse probabilistic graphical models in the non-Gaussian setting
<https://papers.nips.cc/paper/6830-beyond-normality-learning-sparse-probabilistic-graphical-models-in-the-non-gaussian-setting>
AUTHORS: Rebecca Morrison, Ricardo Baptista, Youssef Marzouk
HIGHLIGHT: We present an algorithm to identify sparse dependence structure in continuous and non-Gaussian probability distributions, given a corresponding set of data.

226, TITLE: An inner-loop free solution to inverse problems using deep neural networks
<https://papers.nips.cc/paper/6831-an-inner-loop-free-solution-to-inverse-problems-using-deep-neural-networks>
AUTHORS: Kai Fan, Qi Wei, Lawrence Carin, Katherine A. Heller
HIGHLIGHT: We propose a new method that uses deep learning techniques to accelerate the popular alternating direction method of multipliers (ADMM) solution for inverse problems.

227, TITLE: OnACID: Online Analysis of Calcium Imaging Data in Real Time
<https://papers.nips.cc/paper/6832-onacid-online-analysis-of-calcium-imaging-data-in-real-time>
AUTHORS: Andrea Giovannucci, Johannes Friedrich, Matt Kaufman, Anne Churchland, Dmitri Chklovskii, Liam Paninski, Eftychios A. Pnevmatikakis
HIGHLIGHT: Here we introduce OnACID, an Online framework for the Analysis of streaming Calcium Imaging Data, including i) motion artifact correction, ii) neuronal source extraction, and iii) activity denoising and deconvolution.

228, TITLE: Collaborative PAC Learning
<https://papers.nips.cc/paper/6833-collaborative-pac-learning>
AUTHORS: Avrim Blum, Nika Haghtalab, Ariel D. Procaccia, Mingda Qiao
HIGHLIGHT: We introduce a collaborative PAC learning model, in which k players attempt to learn the same underlying concept.

229, TITLE: Fast Black-box Variational Inference through Stochastic Trust-Region Optimization

- <https://papers.nips.cc/paper/6834-fast-black-box-variational-inference-through-stochastic-trust-region-optimization>
AUTHORS: Jeffrey Regier, Michael I. Jordan, Jon McAuliffe
HIGHLIGHT: We introduce TrustVI, a fast second-order algorithm for black-box variational inference based on trust-region optimization and the reparameterization trick.
- 230, TITLE: Scalable Demand-Aware Recommendation
<https://papers.nips.cc/paper/6835-scalable-demand-aware-recommendation>
AUTHORS: Jinfeng Yi, Cho-Jui Hsieh, Kush R. Varshney, Lijun Zhang, Yao Li
HIGHLIGHT: We further relax this problem and propose a highly scalable alternating minimization approach with which we can solve problems with millions of users and millions of items in a single thread.
- 231, TITLE: SGD Learns the Conjugate Kernel Class of the Network
<https://papers.nips.cc/paper/6836-sgd-learns-the-conjugate-kernel-class-of-the-network>
AUTHORS: Amit Daniely
HIGHLIGHT: We show that the standard stochastic gradient decent (SGD) algorithm is guaranteed to learn, in polynomial time, a function that is competitive with the best function in the conjugate kernel space of the network, as defined in Daniely, Frostig and Singer.
- 232, TITLE: Noise-Tolerant Interactive Learning Using Pairwise Comparisons
<https://papers.nips.cc/paper/6837-noise-tolerant-interactive-learning-using-pairwise-comparisons>
AUTHORS: Yichong Xu, Hongyang Zhang, Kyle Miller, Aarti Singh, Artur Dubrawski
HIGHLIGHT: In this paper, we attempt to characterize how the access to an easier comparison oracle helps in improving the label and total query complexity.
- 233, TITLE: Analyzing Hidden Representations in End-to-End Automatic Speech Recognition Systems
<https://papers.nips.cc/paper/6838-analyzing-hidden-representations-in-end-to-end-automatic-speech-recognition-systems>
AUTHORS: Yonatan Belinkov, James Glass
HIGHLIGHT: In this work, we analyze the speech representations learned by a deep end-to-end model that is based on convolutional and recurrent layers, and trained with a connectionist temporal classification (CTC) loss.
- 234, TITLE: Generative Local Metric Learning for Kernel Regression
<https://papers.nips.cc/paper/6839-generative-local-metric-learning-for-kernel-regression>
AUTHORS: Yung-Kyun Noh, Masashi Sugiyama, Kee-Eung Kim, Frank Park, Daniel D. Lee
HIGHLIGHT: We propose a method for efficiently learning a good metric function based upon analyzing the performance of the NW estimator for Gaussian-distributed data.
- 235, TITLE: Information Theoretic Properties of Markov Random Fields, and their Algorithmic Applications
<https://papers.nips.cc/paper/6840-information-theoretic-properties-of-markov-random-fields-and-their-algorithmic-applications>
AUTHORS: Linus Hamilton, Frederic Koehler, Ankur Moitra
HIGHLIGHT: Here we take a more conceptual approach to proving lower bounds on the mutual information.
- 236, TITLE: Fitting Low-Rank Tensors in Constant Time
<https://papers.nips.cc/paper/6841-fitting-low-rank-tensors-in-constant-time>
AUTHORS: Kohei Hayashi, Yuichi Yoshida
HIGHLIGHT: In this paper, we develop an algorithm that approximates the residual error of Tucker decomposition, one of the most popular tensor decomposition methods, with a provable guarantee.
- 237, TITLE: Deep Supervised Discrete Hashing
<https://papers.nips.cc/paper/6842-deep-supervised-discrete-hashing>
AUTHORS: Qi Li, Zhenan Sun, Ran He, Tieniu Tan
HIGHLIGHT: In this paper, we develop a deep supervised discrete hashing algorithm based on the assumption that the learned binary codes should be ideal for classification.
- 238, TITLE: Using Options and Covariance Testing for Long Horizon Off-Policy Policy Evaluation
<https://papers.nips.cc/paper/6843-using-options-and-covariance-testing-for-long-horizon-off-policy-policy-evaluation>
AUTHORS: Zhaohan Guo, Philip S. Thomas, Emma Brunskill
HIGHLIGHT: We propose using policies over temporally extended actions, called options, and show that combining these policies with importance sampling can significantly improve performance for long-horizon problems.
- 239, TITLE: How regularization affects the critical points in linear networks

<https://papers.nips.cc/paper/6844-how-regularization-affects-the-critical-points-in-linear-networks>
AUTHORS: Amirhossein Taghvaei, Jin W. Kim, Prashant Mehta
HIGHLIGHT: An optimal control model is introduced for this purpose and a learning algorithm (backprop with weight decay) derived for the same using the Hamilton's formulation of optimal control.

240, TITLE: Fisher GAN
<https://papers.nips.cc/paper/6845-fisher-gan>
AUTHORS: Youssef Mroueh, Tom Sercu
HIGHLIGHT: In this paper we introduce Fisher GAN that fits within the Integral Probability Metrics (IPM) framework for training GANs.

241, TITLE: Information-theoretic analysis of generalization capability of learning algorithms
<https://papers.nips.cc/paper/6846-information-theoretic-analysis-of-generalization-capability-of-learning-algorithms>
AUTHORS: Aolin Xu, Maxim Raginsky
HIGHLIGHT: We propose a number of methods for this purpose, among which are algorithms that regularize the ERM algorithm with relative entropy or with random noise.

242, TITLE: Sparse Approximate Conic Hulls
<https://papers.nips.cc/paper/6847-sparse-approximate-conic-hulls>
AUTHORS: Greg Van Buskirk, Benjamin Raichel, Nicholas Ruoizzi
HIGHLIGHT: We consider the problem of computing a restricted nonnegative matrix factorization (NMF) of an $m \times n$ matrix X . Specifically, we seek a factorization $X \approx BC$, where the k columns of B are a subset of those from X and $C \in \mathbb{R}_{\geq 0}^{k \times n}$.

243, TITLE: Rigorous Dynamics and Consistent Estimation in Arbitrarily Conditioned Linear Systems
<https://papers.nips.cc/paper/6848-rigorous-dynamics-and-consistent-estimation-in-arbitrarily-conditioned-linear-systems>
AUTHORS: Alyson K. Fletcher, Mojtaba Sahraee-Ardakan, Sundeep Rangan, Philip Schniter
HIGHLIGHT: The paper thus provides a computationally simple method with provable guarantees of optimality and consistency over a large class of linear inverse problems.

244, TITLE: Toward Goal-Driven Neural Network Models for the Rodent Whisker-Trigeminal System
<https://papers.nips.cc/paper/6849-toward-goal-driven-neural-network-models-for-the-rodent-whisker-trigeminal-system>
AUTHORS: Chengxu Zhuang, Jonas Kubilius, Mitra JZ Hartmann, Daniel L. Yamins
HIGHLIGHT: In the present work, we take a goal-driven deep neural network (DNN) approach to modeling these computations.

245, TITLE: Accuracy First: Selecting a Differential Privacy Level for Accuracy Constrained ERM
<https://papers.nips.cc/paper/6850-accuracy-first-selecting-a-differential-privacy-level-for-accuracy-constrained-erm>
AUTHORS: Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, Steven Z. Wu
HIGHLIGHT: We propose a general “noise reduction” framework that can apply to a variety of private empirical risk minimization (ERM) algorithms, using them to “search” the space of privacy levels to find the empirically strongest one that meets the accuracy constraint, and incurring only logarithmic overhead in the number of privacy levels searched.

246, TITLE: EX2: Exploration with Exemplar Models for Deep Reinforcement Learning
<https://papers.nips.cc/paper/6851-ex2-exploration-with-exemplar-models-for-deep-reinforcement-learning>
AUTHORS: Justin Fu, John Co-Reyes, Sergey Levine
HIGHLIGHT: We propose a novelty detection algorithm for exploration that is based entirely on discriminatively trained exemplar models, where classifiers are trained to discriminate each visited state against all others.

247, TITLE: Multitask Spectral Learning of Weighted Automata
<https://papers.nips.cc/paper/6852-multitask-spectral-learning-of-weighted-automata>
AUTHORS: Guillaume Rabusseau, Borja Balle, Joelle Pineau
HIGHLIGHT: We consider the problem of estimating multiple related functions computed by weighted automata-(WFA).

248, TITLE: Multi-way Interacting Regression via Factorization Machines
<https://papers.nips.cc/paper/6853-multi-way-interacting-regression-via-factorization-machines>
AUTHORS: Mikhail Yurochkin, XuanLong Nguyen, Nikolaos Vasiloglou
HIGHLIGHT: We propose a Bayesian regression method that accounts for multi-way interactions of arbitrary orders among the predictor variables.

- 249, TITLE: Predicting Organic Reaction Outcomes with Weisfeiler-Lehman Network
<https://papers.nips.cc/paper/6854-predicting-organic-reaction-outcomes-with-weisfeiler-lehman-network>
AUTHORS: Wengong Jin, Connor Coley, Regina Barzilay, Tommi Jaakkola
HIGHLIGHT: In this paper, we propose a template-free approach to efficiently explore the space of product molecules by first pinpointing the reaction center -- the set of nodes and edges where graph edits occur.
- 250, TITLE: Practical Data-Dependent Metric Compression with Provable Guarantees
<https://papers.nips.cc/paper/6855-practical-data-dependent-metric-compression-with-provable-guarantees>
AUTHORS: Piotr Indyk, Ilya Razenshteyn, Tal Wagner
HIGHLIGHT: We introduce a new distance-preserving compact representation of multi-dimensional point-sets.
- 251, TITLE: REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models
<https://papers.nips.cc/paper/6856-rebar-low-variance-unbiased-gradient-estimates-for-discrete-latent-variable-models>
AUTHORS: George Tucker, Andriy Mnih, Chris J. Maddison, John Lawson, Jascha Sohl-Dickstein
HIGHLIGHT: In this work, we combine the two approaches through a novel control variate that produces low-variance, unbiased gradient estimates.
- 252, TITLE: Nonlinear random matrix theory for deep learning
<https://papers.nips.cc/paper/6857-nonlinear-random-matrix-theory-for-deep-learning>
AUTHORS: Jeffrey Pennington, Pratik Worah
HIGHLIGHT: In this work, we open the door for direct applications of random matrix theory to deep learning by demonstrating that the pointwise nonlinearities typically applied in neural networks can be incorporated into a standard method of proof in random matrix theory known as the moments method.
- 253, TITLE: Parallel Streaming Wasserstein Barycenters
<https://papers.nips.cc/paper/6858-parallel-streaming-wasserstein-barycenters>
AUTHORS: Matthew Staib, Sebastian Claiici, Justin M. Solomon, Stefanie Jegelka
HIGHLIGHT: Improving on this situation, we present a scalable, communication-efficient, parallel algorithm for computing the Wasserstein barycenter of arbitrary distributions.
- 254, TITLE: ELF: An Extensive, Lightweight and Flexible Research Platform for Real-time Strategy Games
<https://papers.nips.cc/paper/6859-elf-an-extensive-lightweight-and-flexible-research-platform-for-real-time-strategy-games>
AUTHORS: Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, C. Lawrence Zitnick
HIGHLIGHT: In this paper, we propose ELF, an Extensive, Lightweight and Flexible platform for fundamental reinforcement learning research.
- 255, TITLE: Dual Discriminator Generative Adversarial Nets
<https://papers.nips.cc/paper/6860-dual-discriminator-generative-adversarial-nets>
AUTHORS: Tu Nguyen, Trung Le, Hung Vu, Dinh Phung
HIGHLIGHT: We propose in this paper a novel approach to tackle the problem of mode collapse encountered in generative adversarial network (GAN).
- 256, TITLE: Dynamic Revenue Sharing
<https://papers.nips.cc/paper/6861-dynamic-revenue-sharing>
AUTHORS: Santiago Balseiro, Max Lin, Vahab Mirrokni, Renato Leme, IIS Song Zuo
HIGHLIGHT: In this paper, we characterize the optimal revenue sharing scheme that satisfies both constraints in expectation.
- 257, TITLE: Decomposition-Invariant Conditional Gradient for General Polytopes with Line Search
<https://papers.nips.cc/paper/6862-decomposition-invariant-conditional-gradient-for-general-polytopes-with-line-search>
AUTHORS: Mohammad Ali Bashiri, Xinhua Zhang
HIGHLIGHT: In this paper, we show that by employing an away-step update, similar rates can be generalized to arbitrary polytopes with strong empirical performance.
- 258, TITLE: VAIN: Attentional Multi-agent Predictive Modeling
<https://papers.nips.cc/paper/6863-vain-attentional-multi-agent-predictive-modeling>
AUTHORS: Yedid Hoshen
HIGHLIGHT: In this paper we introduce VAIN, a novel attentional architecture for multi-agent predictive modeling that scales linearly with the number of agents.
- 259, TITLE: An Empirical Bayes Approach to Optimizing Machine Learning Algorithms

- <https://papers.nips.cc/paper/6864-an-empirical-bayes-approach-to-optimizing-machine-learning-algorithms>
AUTHORS: James McInerney
HIGHLIGHT: In this paper, we consider an alternative approach that uses more samples from the hyperparameter selection procedure to average over the uncertainty in model hyperparameters.
- 260, TITLE: Differentially Private Empirical Risk Minimization Revisited: Faster and More General
<https://papers.nips.cc/paper/6865-differentially-private-empirical-risk-minimization-revisited-faster-and-more-general>
AUTHORS: Di Wang, Minwei Ye, Jinhui Xu
HIGHLIGHT: In this paper we study differentially private Empirical Risk Minimization(ERM) in different settings.
- 261, TITLE: Variational Inference via χ Upper Bound Minimization
<https://papers.nips.cc/paper/6866-variational-inference-via-chi-upper-bound-minimization>
AUTHORS: Adji Bousso Dieng, Dustin Tran, Rajesh Ranganath, John Paisley, David Blei
HIGHLIGHT: In this paper we propose CHIVI, a black-box variational inference algorithm that minimizes $D_{\chi}(p \parallel q)$, the χ -divergence from p to q .
- 262, TITLE: On Quadratic Convergence of DC Proximal Newton Algorithm in Nonconvex Sparse Learning
<https://papers.nips.cc/paper/6867-on-quadratic-convergence-of-dc-proximal-newton-algorithm-in-nonconvex-sparse-learning>
AUTHORS: Xingguo Li, Lin Yang, Jason Ge, Jarvis Haupt, Tong Zhang, Tuo Zhao
HIGHLIGHT: We propose a DC proximal Newton algorithm for solving nonconvex regularized sparse learning problems in high dimensions.
- 263, TITLE: #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning
<https://papers.nips.cc/paper/6868-exploration-a-study-of-count-based-exploration-for-deep-reinforcement-learning>
AUTHORS: Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, OpenAI Xi Chen, Yan Duan, John Schulman, Filip DeTurck, Pieter Abbeel
HIGHLIGHT: In this work, we describe a surprising finding: a simple generalization of the classic count-based approach can reach near state-of-the-art performance on various high-dimensional and/or continuous deep RL benchmarks.
- 264, TITLE: An Empirical Study on The Properties of Random Bases for Kernel Methods
<https://papers.nips.cc/paper/6869-an-empirical-study-on-the-properties-of-random-bases-for-kernel-methods>
AUTHORS: Maximilian Alber, Pieter-Jan Kindermans, Kristof Sch?tt, Klaus-Robert M?ller, Fei Sha
HIGHLIGHT: In this work, we contrast random features of approximated kernel machines with learned features of neural networks.
- 265, TITLE: Bridging the Gap Between Value and Policy Based Reinforcement Learning
<https://papers.nips.cc/paper/6870-bridging-the-gap-between-value-and-policy-based-reinforcement-learning>
AUTHORS: Ofir Nachum, Mohammad Norouzi, Kelvin Xu, Dale Schuurmans
HIGHLIGHT: From this observation, we develop a new RL algorithm, Path Consistency Learning (PCL), that minimizes a notion of soft consistency error along multi-step action sequences extracted from both on- and off-policy traces.
- 266, TITLE: Premise Selection for Theorem Proving by Deep Graph Embedding
<https://papers.nips.cc/paper/6871-premise-selection-for-theorem-proving-by-deep-graph-embedding>
AUTHORS: Mingzhe Wang, Yihe Tang, Jian Wang, Jia Deng
HIGHLIGHT: We propose a deep learning-based approach to the problem of premise selection: selecting mathematical statements relevant for proving a given conjecture.
- 267, TITLE: A Bayesian Data Augmentation Approach for Learning Deep Models
<https://papers.nips.cc/paper/6872-a-bayesian-data-augmentation-approach-for-learning-deep-models>
AUTHORS: Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, Ian Reid
HIGHLIGHT: For learning, we introduce a theoretically sound algorithm --- generalised Monte Carlo expectation maximisation, and demonstrate one possible implementation via an extension of the Generative Adversarial Network (GAN).
- 268, TITLE: Principles of Riemannian Geometry in Neural Networks
<https://papers.nips.cc/paper/6873-principles-of-riemannian-geometry-in-neural-networks>
AUTHORS: Michael Hauser, Asok Ray
HIGHLIGHT: This study deals with neural networks in the sense of geometric transformations acting on the coordinate representation of the underlying data manifold which the data is sampled from.
- 269, TITLE: Cold-Start Reinforcement Learning with Softmax Policy Gradient

<https://papers.nips.cc/paper/6874-cold-start-reinforcement-learning-with-softmax-policy-gradient>

AUTHORS: Nan Ding, Radu Soricut

HIGHLIGHT: In this paper, we describe a reinforcement learning method based on a softmax value function that requires neither of these procedures.

270, TITLE: Online Dynamic Programming

<https://papers.nips.cc/paper/6875-online-dynamic-programming>

AUTHORS: Holakou Rahmanian, Manfred K. Warmuth

HIGHLIGHT: We consider the problem of repeatedly solving a variant of the same dynamic programming problem in successive trials.

271, TITLE: Alternating Estimation for Structured High-Dimensional Multi-Response Models

<https://papers.nips.cc/paper/6876-alternating-estimation-for-structured-high-dimensional-multi-response-models>

AUTHORS: Sheng Chen, Arindam Banerjee

HIGHLIGHT: By exploiting the noise correlations among different responses, we propose an alternating estimation (AltEst) procedure to estimate the model parameters based on the generalized Dantzig selector (GDS).

272, TITLE: Convolutional Gaussian Processes

<https://papers.nips.cc/paper/6877-convolutional-gaussian-processes>

AUTHORS: Mark van der Wilk, Carl Edward Rasmussen, James Hensman

HIGHLIGHT: The main contribution of our work is the construction of an inter-domain inducing point approximation that is well-tailored to the convolutional kernel.

273, TITLE: Estimation of the covariance structure of heavy-tailed distributions

<https://papers.nips.cc/paper/6878-estimation-of-the-covariance-structure-of-heavy-tailed-distributions>

AUTHORS: Xiaohan Wei, Stanislav Minsker

HIGHLIGHT: We propose and analyze a new estimator of the covariance matrix that admits strong theoretical guarantees under weak assumptions on the underlying distribution, such as existence of moments of only low order.

274, TITLE: Decomposable Submodular Function Minimization: Discrete and Continuous

<https://papers.nips.cc/paper/6880-decomposable-submodular-function-minimization-discrete-and-continuous>

AUTHORS: Alina Ene, Huy Nguyen, L?szl? A. V?gh

HIGHLIGHT: This paper investigates connections between discrete and continuous approaches for decomposable submodular function minimization.

275, TITLE: Gauging Variational Inference

<https://papers.nips.cc/paper/6881-gauging-variational-inference>

AUTHORS: Sung-Soo Ahn, Michael Chertkov, Jinwoo Shin

HIGHLIGHT: In this paper, we propose two new variational schemes, coined Gauged-MF (G-MF) and Gauged-BP (G-BP), improving MF and BP, respectively.

276, TITLE: Deep Recurrent Neural Network-Based Identification of Precursor microRNAs

<https://papers.nips.cc/paper/6882-deep-recurrent-neural-network-based-identification-of-precursor-micrnas>

AUTHORS: Seunghyun Park, Seonwoo Min, Hyun-Soo Choi, Sungroh Yoon

HIGHLIGHT: To address the limitations of existing approaches, we propose a pre-miRNA identification method that incorporates (1) a deep recurrent neural network (RNN) for automated feature learning and classification, (2) multimodal architecture for seamless integration of prior knowledge (secondary structure), (3) an attention mechanism for improving long-term dependence modeling, and (4) an RNN-based class activation mapping for highlighting the learned representations that can contrast pre-miRNAs and non-pre-miRNAs.

277, TITLE: Robust Estimation of Neural Signals in Calcium Imaging

<https://papers.nips.cc/paper/6883-robust-estimation-of-neural-signals-in-calcium-imaging>

AUTHORS: Hakan Inan, Murat A. Erdogdu, Mark Schnitzer

HIGHLIGHT: In this work, we proceed in a new direction and propose to extract cells and their activity using robust statistical estimation.

278, TITLE: State Aware Imitation Learning

<https://papers.nips.cc/paper/6884-state-aware-imitation-learning>

AUTHORS: Yannick Schroecker, Charles L. Isbell

HIGHLIGHT: In this paper, we introduce State Aware Imitation Learning (SAIL), an imitation learning algorithm that allows an agent to learn how to remain in states where it can confidently take the correct action and how to recover if it is lead astray.

- 279, TITLE: Beyond Parity: Fairness Objectives for Collaborative Filtering
<https://papers.nips.cc/paper/6885-beyond-parity-fairness-objectives-for-collaborative-filtering>
AUTHORS: Sirui Yao, Bert Huang
HIGHLIGHT: We identify the insufficiency of existing fairness metrics and propose four new metrics that address different forms of unfairness.
- 280, TITLE: A PAC-Bayesian Analysis of Randomized Learning with Application to Stochastic Gradient Descent
<https://papers.nips.cc/paper/6886-a-pac-bayesian-analysis-of-randomized-learning-with-application-to-stochastic-gradient-descent>
AUTHORS: Ben London
HIGHLIGHT: We analyze this algorithm in the context of our generalization bounds and evaluate it on a benchmark dataset.
- 281, TITLE: Fully Decentralized Policies for Multi-Agent Systems: An Information Theoretic Approach
<https://papers.nips.cc/paper/6887-fully-decentralized-policies-for-multi-agent-systems-an-information-theoretic-approach>
AUTHORS: Roel Dobbe, David Fridovich-Keil, Claire Tomlin
HIGHLIGHT: Our main contribution is an information theoretic framework based on rate distortion theory which facilitates analysis of how well the resulting fully decentralized policies are able to reconstruct the optimal solution.
- 282, TITLE: Model-Powered Conditional Independence Test
<https://papers.nips.cc/paper/6888-model-powered-conditional-independence-test>
AUTHORS: Rajat Sen, Ananda Theertha Suresh, Karthikeyan Shanmugam, Alexandros G. Dimakis, Sanjay Shakkottai
HIGHLIGHT: To tackle this problem we propose a novel nearest neighbor bootstrap procedure and theoretically show that our generated samples are indeed close to \mathcal{P}^{CI} in terms of total variational distance.
- 283, TITLE: Deep Voice 2: Multi-Speaker Neural Text-to-Speech
<https://papers.nips.cc/paper/6889-deep-voice-2-multi-speaker-neural-text-to-speech>
AUTHORS: Andrew Gibiansky, Sercan Arik, Gregory Diamos, John Miller, Kainan Peng, Wei Ping, Jonathan Raiman, Yanqi Zhou
HIGHLIGHT: We introduce a technique for augmenting neural text-to-speech (TTS) with low-dimensional trainable speaker embeddings to generate different voices from a single model.
- 284, TITLE: Variance-based Regularization with Convex Objectives
<https://papers.nips.cc/paper/6890-variance-based-regularization-with-convex-objectives>
AUTHORS: Hongseok Namkoong, John C. Duchi
HIGHLIGHT: We develop an approach to risk minimization and stochastic optimization that provides a convex surrogate for variance, allowing near-optimal and computationally efficient trading between approximation and estimation error.
- 285, TITLE: Deep Lattice Networks and Partial Monotonic Functions
<https://papers.nips.cc/paper/6891-deep-lattice-networks-and-partial-monotonic-functions>
AUTHORS: Seungil You, David Ding, Kevin Canini, Jan Pfeifer, Maya Gupta
HIGHLIGHT: We propose learning deep models that are monotonic with respect to a user-specified set of inputs by alternating layers of linear embeddings, ensembles of lattices, and calibrators (piecewise linear functions), with appropriate constraints for monotonicity, and jointly training the resulting network.
- 286, TITLE: Continual Learning with Deep Generative Replay
<https://papers.nips.cc/paper/6892-continual-learning-with-deep-generative-replay>
AUTHORS: Hanul Shin, Jung Kwon Lee, Jaehong Kim, Jiwon Kim
HIGHLIGHT: Inspired by the generative nature of the hippocampus as a short-term memory system in primate brain, we propose the Deep Generative Replay, a novel framework with a cooperative dual model architecture consisting of a deep generative model (“generator”) and a task solving model (“solver”).
- 287, TITLE: AIDE: An algorithm for measuring the accuracy of probabilistic inference algorithms
<https://papers.nips.cc/paper/6893-aide-an-algorithm-for-measuring-the-accuracy-of-probabilistic-inference-algorithms>
AUTHORS: Marco Cusumano-Towner, Vikash K. Mansinghka
HIGHLIGHT: This paper introduces the auxiliary inference divergence estimator (AIDE), an algorithm for measuring the accuracy of approximate inference algorithms.
- 288, TITLE: Learning Causal Structures Using Regression Invariance
<https://papers.nips.cc/paper/6894-learning-causal-structures-using-regression-invariance>
AUTHORS: AmirEmad Ghassami, Saber Salehkaleybar, Negar Kiyavash, Kun Zhang

HIGHLIGHT: We introduce the idea of using the invariance of the functional relations of the variables to their causes across a set of environments for structure learning.

289, TITLE: Online Influence Maximization under Independent Cascade Model with Semi-Bandit Feedback
<https://papers.nips.cc/paper/6895-online-influence-maximization-under-independent-cascade-model-with-semi-bandit-feedback>
AUTHORS: Zheng Wen, Branislav Kveton, Michal Valko, Sharan Vaswani
HIGHLIGHT: Under a stochastic semi-bandit feedback, we propose and analyze IMLinUCB, a computationally efficient UCB-based algorithm.

290, TITLE: Near Minimax Optimal Players for the Finite-Time 3-Expert Prediction Problem
<https://papers.nips.cc/paper/6896-near-minimax-optimal-players-for-the-finite-time-3-expert-prediction-problem>
AUTHORS: Yasin Abbasi, Peter L. Bartlett, Victor Gabillon
HIGHLIGHT: We study minimax strategies for the online prediction problem with expert advice.

291, TITLE: Reinforcement Learning under Model Mismatch
<https://papers.nips.cc/paper/6897-reinforcement-learning-under-model-mismatch>
AUTHORS: Aurko Roy, Huan Xu, Sebastian Pokutta
HIGHLIGHT: We address this problem by extending the framework of robust MDPs to the model-free Reinforcement Learning setting, where we do not have access to the model parameters, but can only sample states from it.

292, TITLE: Hierarchical Attentive Recurrent Tracking
<https://papers.nips.cc/paper/6898-hierarchical-attentive-recurrent-tracking>
AUTHORS: Adam Kosiosek, Alex Bewley, Ingmar Posner
HIGHLIGHT: Inspired by how the human visual cortex employs spatial attention and separate "where" and "what" processing pathways to actively suppress irrelevant visual features, this work develops a hierarchical attentive recurrent model for single object tracking in videos.

293, TITLE: Tomography of the London Underground: a Scalable Model for Origin-Destination Data
<https://papers.nips.cc/paper/6899-tomography-of-the-london-underground-a-scalable-model-for-origin-destination-data>
AUTHORS: Nicolò Colombo, Ricardo Silva, Soong Moon Kang
HIGHLIGHT: The corresponding likelihood function is intractable for medium or large-scale networks and we propose two distinct strategies, namely the exact maximum-likelihood inference of an approximate but tractable model and the variational inference of the original intractable model.

294, TITLE: Rotting Bandits
<https://papers.nips.cc/paper/6900-rotting-bandits>
AUTHORS: Nir Levine, Koby Crammer, Shie Mannor
HIGHLIGHT: We present algorithms, accompanied by simulations, and derive theoretical guarantees.

295, TITLE: Unbiased estimates for linear regression via volume sampling
<https://papers.nips.cc/paper/6901-unbiased-estimates-for-linear-regression-via-volume-sampling>
AUTHORS: Michal Dereziński, Manfred K. Warmuth
HIGHLIGHT: Pseudo inverse plays an important part in solving the linear least squares problem, where we try to predict a label for each column of XX^T .

296, TITLE: Approximation Bounds for Hierarchical Clustering: Average Linkage, Bisecting K-means, and Local Search
<https://papers.nips.cc/paper/6902-approximation-bounds-for-hierarchical-clustering-average-linkage-bisecting-k-means-and-local-search>
AUTHORS: Benjamin Moseley, Joshua Wang
HIGHLIGHT: The goal of this paper is to give an analytic framework to better understand observations seen in practice.

297, TITLE: Adaptive Accelerated Gradient Converging Method under Hölderian Error Bound Condition
<https://papers.nips.cc/paper/6903-adaptive-accelerated-gradient-converging-method-under-holderian-error-bound-condition>
AUTHORS: Mingrui Liu, Tianbao Yang
HIGHLIGHT: To the best of our knowledge, these improved results are first shown in this work.

298, TITLE: Stein Variational Gradient Descent as Gradient Flow
<https://papers.nips.cc/paper/6904-stein-variational-gradient-descent-as-gradient-flow>
AUTHORS: Qiang Liu
HIGHLIGHT: This paper develops the first theoretical analysis on SVGD.

- 299, TITLE: Partial Hard Thresholding: Towards A Principled Analysis of Support Recovery
<https://papers.nips.cc/paper/6905-partial-hard-thresholding-towards-a-principled-analysis-of-support-recovery>
AUTHORS: Jie Shen, Ping Li
HIGHLIGHT: In this paper, we present a principled analysis on the support recovery performance for a family of hard thresholding algorithms.
- 300, TITLE: Shallow Updates for Deep Reinforcement Learning
<https://papers.nips.cc/paper/6906-shallow-updates-for-deep-reinforcement-learning>
AUTHORS: Nir Levine, Tom Zahavy, Daniel J. Mankowitz, Aviv Tamar, Shie Mannor
HIGHLIGHT: In this work we propose a hybrid approach -- the Least Squares Deep Q-Network (LS-DQN), which combines rich feature representations learned by a DRL algorithm with the stability of a linear least squares method.
- 301, TITLE: LightGBM: A Highly Efficient Gradient Boosting Decision Tree
<https://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree>
AUTHORS: Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, Tie-Yan Liu
HIGHLIGHT: To tackle this problem, we propose two novel techniques: `\emph{Gradient-based One-Side Sampling}` (GOSS) and `\emph{Exclusive Feature Bundling}` (EFB).
- 302, TITLE: Adversarial Ranking for Language Generation
<https://papers.nips.cc/paper/6908-adversarial-ranking-for-language-generation>
AUTHORS: Kevin Lin, Dianqi Li, Xiaodong He, Zhengyou Zhang, Ming-ting Sun
HIGHLIGHT: In this paper, we propose a novel generative adversarial network, RankGAN, for generating high-quality language descriptions.
- 303, TITLE: Regret Minimization in MDPs with Options without Prior Knowledge
<https://papers.nips.cc/paper/6909-regret-minimization-in-mdps-with-options-without-prior-knowledge>
AUTHORS: Ronan Fruit, Matteo Pirota, Alessandro Lazaric, Emma Brunskill
HIGHLIGHT: In this paper, we remove this limitation by combining the SMDP view together with the inner Markov structure of options into a novel algorithm whose regret performance matches UCRL-SMDP's up to an additive regret term.
- 304, TITLE: Net-Trim: Convex Pruning of Deep Neural Networks with Performance Guarantee
<https://papers.nips.cc/paper/6910-net-trim-convex-pruning-of-deep-neural-networks-with-performance-guarantee>
AUTHORS: Alireza Aghasi, Afshin Abdi, Nam Nguyen, Justin Romberg
HIGHLIGHT: We introduce and analyze a new technique for model reduction for deep neural networks.
- 305, TITLE: Graph Matching via Multiplicative Update Algorithm
<https://papers.nips.cc/paper/6911-graph-matching-via-multiplicative-update-algorithm>
AUTHORS: Bo Jiang, Jin Tang, Chris Ding, Yihong Gong, Bin Luo
HIGHLIGHT: In this paper, we present a new algorithm, called Multiplicative Update Graph Matching (MPGM), that develops a multiplicative update technique to solve the QP matching problem.
- 306, TITLE: Dynamic Importance Sampling for Anytime Bounds of the Partition Function
<https://papers.nips.cc/paper/6912-dynamic-importance-sampling-for-anytime-bounds-of-the-partition-function>
AUTHORS: Qi Lou, Rina Dechter, Alexander T. Ihler
HIGHLIGHT: In this paper, we propose a dynamic importance sampling scheme that provides anytime finite-sample bounds for the partition function.
- 307, TITLE: Is the Bellman residual a bad proxy?
<https://papers.nips.cc/paper/6913-is-the-bellman-residual-a-bad-proxy>
AUTHORS: Matthieu Geist, Bilal Piot, Olivier Pietquin
HIGHLIGHT: This paper aims at theoretically and empirically comparing two standard optimization criteria for Reinforcement Learning: i) maximization of the mean value and ii) minimization of the Bellman residual.
- 308, TITLE: Generalization Properties of Learning with Random Features
<https://papers.nips.cc/paper/6914-generalization-properties-of-learning-with-random-features>
AUTHORS: Alessandro Rudi, Lorenzo Rosasco
HIGHLIGHT: We study the generalization properties of ridge regression with random features in the statistical learning framework.

- 309, TITLE: Differentially private Bayesian learning on distributed data
<https://papers.nips.cc/paper/6915-differentially-private-bayesian-learning-on-distributed-data>
AUTHORS: Mikko Heikkilä, Emil Lagerspetz, Samuel Kaski, Kana Shimizu, Sasu Tarkoma, Antti Honkela
HIGHLIGHT: We propose a learning strategy based on a secure multi-party sum function for aggregating summaries from data holders and the Gaussian mechanism for DP.
- 310, TITLE: Learning to Compose Domain-Specific Transformations for Data Augmentation
<https://papers.nips.cc/paper/6916-learning-to-compose-domain-specific-transformations-for-data-augmentation>
AUTHORS: Alexander J. Ratner, Henry Ehrenberg, Zeshan Hussain, Jared Dunnmon, Christopher R?
HIGHLIGHT: We propose a method for automating this process by learning a generative sequence model over user-specified transformation functions using a generative adversarial approach.
- 311, TITLE: Wasserstein Learning of Deep Generative Point Process Models
<https://papers.nips.cc/paper/6917-wasserstein-learning-of-deep-generative-point-process-models>
AUTHORS: Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, Hongyuan Zha
HIGHLIGHT: In this paper, we propose an intensity-free approach for point processes modeling that transforms nuisance processes to a target one.
- 312, TITLE: Ensemble Sampling
<https://papers.nips.cc/paper/6918-ensemble-sampling>
AUTHORS: Xiuyuan Lu, Benjamin Van Roy
HIGHLIGHT: This paper develops ensemble sampling, which aims to approximate Thompson sampling while maintaining tractability even in the face of complex models such as neural networks.
- 313, TITLE: Language Modeling with Recurrent Highway Hypernetworks
<https://papers.nips.cc/paper/6919-language-modeling-with-recurrent-highway-hypernetworks>
AUTHORS: Joseph Suarez
HIGHLIGHT: We present extensive experimental and theoretical support for the efficacy of recurrent highway networks (RHNs) and recurrent hypernetworks complementary to the original works.
- 314, TITLE: Adaptive SVRG Methods under Error Bound Conditions with Unknown Growth Parameter
<https://papers.nips.cc/paper/6920-adaptive-svrg-methods-under-error-bound-conditions-with-unknown-growth-parameter>
AUTHORS: Yi Xu, Qihang Lin, Tianbao Yang
HIGHLIGHT: To address this issue, we propose novel SVRG methods that automatically search for this unknown parameter on the fly of optimization while still obtain almost the same convergence rate as when this parameter is known.
- 315, TITLE: Bayesian Compression for Deep Learning
<https://papers.nips.cc/paper/6921-bayesian-compression-for-deep-learning>
AUTHORS: Christos Louizos, Karen Ullrich, Max Welling
HIGHLIGHT: In this work, we argue that the most principled and effective way to attack this problem is by adopting a Bayesian point of view, where through sparsity inducing priors we prune large parts of the network.
- 316, TITLE: Streaming Sparse Gaussian Process Approximations
<https://papers.nips.cc/paper/6922-streaming-sparse-gaussian-process-approximations>
AUTHORS: Thang D. Bui, Cuong Nguyen, Richard E. Turner
HIGHLIGHT: This paper develops a new principled framework for deploying Gaussian process probabilistic models in the streaming setting, providing methods for learning hyperparameters and optimising pseudo-input locations.
- 317, TITLE: VEEGAN: Reducing Mode Collapse in GANs using Implicit Variational Learning
<https://papers.nips.cc/paper/6923-veegan-reducing-mode-collapse-in-gans-using-implicit-variational-learning>
AUTHORS: Akash Srivastava, Lazar Valkov, Chris Russell, Michael U. Gutmann, Charles Sutton
HIGHLIGHT: To address this, we introduce VEEGAN, which features a reconstructor network, reversing the action of the generator by mapping from data to noise.
- 318, TITLE: Sparse Embedded k -Means Clustering
<https://papers.nips.cc/paper/6924-sparse-embedded-k-means-clustering>
AUTHORS: Weiwei Liu, Xiaobo Shen, Ivor Tsang
HIGHLIGHT: Our empirical studies corroborate our theoretical findings, and demonstrate that our approach is able to significantly accelerate k -means clustering, while achieving satisfactory clustering performance.

- 319, TITLE: Dynamic-Depth Context Tree Weighting
<https://papers.nips.cc/paper/6925-dynamic-depth-context-tree-weighting>
AUTHORS: Joao V. Messias, Shimon Whiteson
HIGHLIGHT: This paper proposes dynamic-depth context tree weighting (D2-CTW), a model-learning method that addresses these limitations.
- 320, TITLE: A Regularized Framework for Sparse and Structured Neural Attention
<https://papers.nips.cc/paper/6926-a-regularized-framework-for-sparse-and-structured-neural-attention>
AUTHORS: Vlad Niculae, Mathieu Blondel
HIGHLIGHT: We propose in this paper a new framework for sparse and structured attention, building upon a smoothed max operator.
- 321, TITLE: Multi-output Polynomial Networks and Factorization Machines
<https://papers.nips.cc/paper/6927-multi-output-polynomial-networks-and-factorization-machines>
AUTHORS: Mathieu Blondel, Vlad Niculae, Takuma Otsuka, Naonori Ueda
HIGHLIGHT: We extend these models to the multi-output setting, i.e., for learning vector-valued functions, with application to multi-class or multi-task problems.
- 322, TITLE: Clustering Billions of Reads for DNA Data Storage
<https://papers.nips.cc/paper/6928-clustering-billions-of-reads-for-dna-data-storage>
AUTHORS: Cyrus Rashtchian, Konstantin Makarychev, Miklos Racz, Siena Ang, Djordje Jevdjic, Sergey Yekhanin, Luis Ceze, Karin Strauss
HIGHLIGHT: To address this issue, we present a novel distributed algorithm for approximately computing the underlying clusters.
- 323, TITLE: Multi-Objective Non-parametric Sequential Prediction
<https://papers.nips.cc/paper/6929-multi-objective-non-parametric-sequential-prediction>
AUTHORS: Guy Uziel, Ran El-Yaniv
HIGHLIGHT: In this paper, we extend the multi-objective framework to the case of stationary and ergodic processes, thus allowing dependencies among observations.
- 324, TITLE: A Universal Analysis of Large-Scale Regularized Least Squares Solutions
<https://papers.nips.cc/paper/6930-a-universal-analysis-of-large-scale-regularized-least-squares-solutions>
AUTHORS: Ashkan Panahi, Babak Hassibi
HIGHLIGHT: In this paper, we confirm these observations theoretically when the distribution is sub-Gaussian.
- 325, TITLE: Deep Sets
<https://papers.nips.cc/paper/6931-deep-sets>
AUTHORS: Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R. Salakhutdinov, Alexander J. Smola
HIGHLIGHT: We study the problem of designing models for machine learning tasks defined on sets.
- 326, TITLE: ExtremeWeather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events
<https://papers.nips.cc/paper/6932-extremeweather-a-large-scale-climate-dataset-for-semi-supervised-detection-localization-and-understanding-of-extreme-weather-events>
AUTHORS: Evan Racah, Christopher Beckham, Tegan Maharaj, Samira Ebrahimi Kahou, Mr. Prabhat, Chris Pal
HIGHLIGHT: We present a multichannel spatiotemporal CNN architecture for semi-supervised bounding box prediction and exploratory data analysis.
We present a dataset, ExtremeWeather, to encourage machine learning research in this area and to help facilitate further work in understanding and mitigating the effects of climate change.
- 327, TITLE: Process-constrained batch Bayesian optimisation
<https://papers.nips.cc/paper/6933-process-constrained-batch-bayesian-optimisation>
AUTHORS: Pratibha Vellanki, Santu Rana, Sunil Gupta, David Rubin, Alessandra Sutti, Thomas Dorin, Murray Height, Paul Sanders, Svetha Venkatesh
HIGHLIGHT: We propose two algorithms, pc-BO(basic) and pc-BO(nested).
- 328, TITLE: Bayesian Inference of Individualized Treatment Effects using Multi-task Gaussian Processes
<https://papers.nips.cc/paper/6934-bayesian-inference-of-individualized-treatment-effects-using-multi-task-gaussian-processes>

- AUTHORS: Ahmed M. Alaa, Mihaela van der Schaar
HIGHLIGHT: Stemming from the potential outcomes model, we propose a novel multi-task learning framework in which factual and counterfactual outcomes are modeled as the outputs of a function in a vector-valued reproducing kernel Hilbert space (vvRKHS).
- 329, TITLE: Spherical convolutions and their application in molecular modelling
<https://papers.nips.cc/paper/6935-spherical-convolutions-and-their-application-in-molecular-modelling>
AUTHORS: Wouter Boomsma, Jes Frellsen
HIGHLIGHT: In this paper, we introduce two strategies for conducting convolutions on the sphere, using either a spherical-polar grid or a grid based on the cubed-sphere representation.
- 330, TITLE: Efficient Optimization for Linear Dynamical Systems with Applications to Clustering and Sparse Coding
<https://papers.nips.cc/paper/6936-efficient-optimization-for-linear-dynamical-systems-with-applications-to-clustering-and-sparse-coding>
AUTHORS: Wenbing Huang, Mehrtaash Harandi, Tong Zhang, Lijie Fan, Fuchun Sun, Junzhou Huang
HIGHLIGHT: In this paper, we propose an efficient projected gradient descent method to minimize a general form of a loss function and demonstrate how clustering and sparse coding with LDSs can be solved by the proposed method efficiently.
- 331, TITLE: On Optimal Generalizability in Parametric Learning
<https://papers.nips.cc/paper/6937-on-optimal-generalizability-in-parametric-learning>
AUTHORS: Ahmad Beirami, Meisam Razaviyayn, Shahin Shahrampour, Vahid Tarokh
HIGHLIGHT: We consider the parametric learning problem, where the objective of the learner is determined by a parametric loss function.
- 332, TITLE: Near Optimal Sketching of Low-Rank Tensor Regression
<https://papers.nips.cc/paper/6938-near-optimal-sketching-of-low-rank-tensor-regression>
AUTHORS: Xingguo Li, Jarvis Haupt, David Woodruff
HIGHLIGHT: We study the least squares regression problem $\min_{\Theta} \|\Theta\|_{\text{F}}^2$ where Θ is a low-rank tensor, defined as $\Theta = \sum_{r=1}^R \theta_r \circ \theta_r$, for vectors θ_r for all $r \in [R]$ and $d \in [D]$.
- 333, TITLE: Tractability in Structured Probability Spaces
<https://papers.nips.cc/paper/6939-tractability-in-structured-probability-spaces>
AUTHORS: Arthur Choi, Yujia Shen, Adnan Darwiche
HIGHLIGHT: In this paper, we study the scalability of such models in the context of representing and learning distributions over routes on a map.
- 334, TITLE: Model-based Bayesian inference of neural activity and connectivity from all-optical interrogation of a neural circuit
<https://papers.nips.cc/paper/6940-model-based-bayesian-inference-of-neural-activity-and-connectivity-from-all-optical-interrogation-of-a-neural-circuit>
AUTHORS: Laurence Aitchison, Lloyd Russell, Adam M. Packer, Jinyao Yan, Philippe Castonguay, Michael Hausser, Srinivas C. Turaga
HIGHLIGHT: We use the framework of variational autoencoders to model spiking activity using discrete latent variables, low-dimensional latent common input, and sparse spike-and-slab generalized linear coupling between neurons.
- 335, TITLE: Gaussian process based nonlinear latent structure discovery in multivariate spike train data
<https://papers.nips.cc/paper/6941-gaussian-process-based-nonlinear-latent-structure-discovery-in-multivariate-spike-train-data>
AUTHORS: Anqi Wu, Nicholas A. Roy, Stephen Keeley, Jonathan W. Pillow
HIGHLIGHT: Here we propose a doubly nonlinear latent variable model that can identify low-dimensional structure underlying apparently high-dimensional spike train data.
- 336, TITLE: Neural system identification for large populations separating ?what? and ?where?
<https://papers.nips.cc/paper/6942-neural-system-identification-for-large-populations-separating-what-and-where>
AUTHORS: David Klindt, Alexander S. Ecker, Thomas Euler, Matthias Bethge
HIGHLIGHT: We propose a CNN architecture with a sparse readout layer factorizing the spatial (where) and feature (what) dimensions.
- 337, TITLE: Certified Defenses for Data Poisoning Attacks
<https://papers.nips.cc/paper/6943-certified-defenses-for-data-poisoning-attacks>
AUTHORS: Jacob Steinhardt, Pang Wei W. Koh, Percy S. Liang

HIGHLIGHT: Empirically, we find that even under a simple defense, the MNIST-1-7 and Dogfish datasets are resilient to attack, while in contrast the IMDB sentiment dataset can be driven from 12% to 23% test error by adding only 3% poisoned data.

338, **TITLE:** Eigen-Distortions of Hierarchical Representations
<https://papers.nips.cc/paper/6944-eigen-distortions-of-hierarchical-representations>
AUTHORS: Alexander Berardino, Valero Laparra, Johannes Ball?, Eero Simoncelli
HIGHLIGHT: We develop a method for comparing hierarchical image representations in terms of their ability to explain perceptual sensitivity in humans.

339, **TITLE:** Limitations on Variance-Reduction and Acceleration Schemes for Finite Sums Optimization
<https://papers.nips.cc/paper/6945-limitations-on-variance-reduction-and-acceleration-schemes-for-finite-sums-optimization>
AUTHORS: Yossi Arjevani
HIGHLIGHT: We study the conditions under which one is able to efficiently apply variance-reduction and acceleration schemes on finite sums problems.

340, **TITLE:** Unsupervised Sequence Classification using Sequential Output Statistics
<https://papers.nips.cc/paper/6946-unsupervised-sequence-classification-using-sequential-output-statistics>
AUTHORS: Yu Liu, Jianshu Chen, Li Deng
HIGHLIGHT: To address the problem, we propose an unsupervised learning cost function and study its properties.

341, **TITLE:** Subset Selection under Noise
<https://papers.nips.cc/paper/6947-subset-selection-under-noise>
AUTHORS: Chao Qian, Jing-Cheng Shi, Yang Yu, Ke Tang, Zhi-Hua Zhou
HIGHLIGHT: We prove that PONSS can achieve a better approximation ratio under some assumption such as i.i.d. noise distribution.

342, **TITLE:** Collecting Telemetry Data Privately
<https://papers.nips.cc/paper/6948-collecting-telemetry-data-privately>
AUTHORS: Bolin Ding, Janardhan Kulkarni, Sergey Yekhanin
HIGHLIGHT: In this paper, we develop new LDP mechanisms geared towards repeated collection of counter data, with formal privacy guarantees even after being executed for an arbitrarily long period of time.

343, **TITLE:** Concrete Dropout
<https://papers.nips.cc/paper/6949-concrete-dropout>
AUTHORS: Yarin Gal, Jiri Hron, Alex Kendall
HIGHLIGHT: We propose a new dropout variant which gives improved performance and better calibrated uncertainties.

344, **TITLE:** Adaptive Batch Size for Safe Policy Gradients
<https://papers.nips.cc/paper/6950-adaptive-batch-size-for-safe-policy-gradients>
AUTHORS: Matteo Papini, Matteo Pirota, Marcello Restelli
HIGHLIGHT: In this paper, we propose a set of methods to jointly optimize the step and the batch sizes that guarantee (with high probability) to improve the policy performance after each update.

345, **TITLE:** A Disentangled Recognition and Nonlinear Dynamics Model for Unsupervised Learning
<https://papers.nips.cc/paper/6951-a-disentangled-recognition-and-nonlinear-dynamics-model-for-unsupervised-learning>
AUTHORS: Marco Fraccaro, Simon Kamronn, Ulrich Paquet, Ole Winther
HIGHLIGHT: We introduce the Kalman variational auto-encoder, a framework for unsupervised learning of sequential data that disentangles two latent representations: an object's representation, coming from a recognition model, and a latent state describing its dynamics.

346, **TITLE:** PASS-GLM: polynomial approximate sufficient statistics for scalable Bayesian GLM inference
<https://papers.nips.cc/paper/6952-pass-glm-polynomial-approximate-sufficient-statistics-for-scalable-bayesian-glm-inference>
AUTHORS: Jonathan Huggins, Ryan P. Adams, Tamara Broderick
HIGHLIGHT: We propose a new approach based on constructing polynomial approximate sufficient statistics for GLMs (PASS-GLM).

347, **TITLE:** Bayesian GAN
<https://papers.nips.cc/paper/6953-bayesian-gan>
AUTHORS: Yunus Saatci, Andrew G. Wilson
HIGHLIGHT: We present a practical Bayesian formulation for unsupervised and semi-supervised learning with GANs.

- 348, TITLE: Off-policy evaluation for slate recommendation
<https://papers.nips.cc/paper/6954-off-policy-evaluation-for-slate-recommendation>
AUTHORS: Adith Swaminathan, Akshay Krishnamurthy, Alekh Agarwal, Miro Dudik, John Langford, Damien Jose, Imed Zitouni
HIGHLIGHT: We build on techniques from combinatorial bandits to introduce a new practical estimator that uses logged data to estimate a policy's performance.
- 349, TITLE: A multi-agent reinforcement learning model of common-pool resource appropriation
<https://papers.nips.cc/paper/6955-a-multi-agent-reinforcement-learning-model-of-common-pool-resource-appropriation>
AUTHORS: Julien P?rolat, Joel Z. Leibo, Vinicius Zambaldi, Charles Beattie, Karl Tuyls, Thore Graepel
HIGHLIGHT: Abstract models of common-pool resource appropriation based on non-cooperative game theory predict that self-interested agents will generally fail to find socially positive equilibria---a phenomenon called the tragedy of the commons.
- 350, TITLE: On the Optimization Landscape of Tensor Decompositions
<https://papers.nips.cc/paper/6956-on-the-optimization-landscape-of-tensor-decompositions>
AUTHORS: Rong Ge, Tengyu Ma
HIGHLIGHT: In this paper, we analyze the optimization landscape of the random over-complete tensor decomposition problem, which has many applications in unsupervised learning, especially in learning latent variable models.
- 351, TITLE: High-Order Attention Models for Visual Question Answering
<https://papers.nips.cc/paper/6957-high-order-attention-models-for-visual-question-answering>
AUTHORS: Idan Schwartz, Alexander Schwing, Tamir Hazan
HIGHLIGHT: In this paper we propose a novel and generally applicable form of attention mechanism that learns high-order correlations between various data modalities.
- 352, TITLE: Sparse convolutional coding for neuronal assembly detection
<https://papers.nips.cc/paper/6958-sparse-convolutional-coding-for-neuronal-assembly-detection>
AUTHORS: Sven Peter, Elke Kirschbaum, Martin Both, Lee Campbell, Brandon Harvey, Conor Heins, Daniel Durstewitz, Ferran Diego, Fred A. Hamprecht
HIGHLIGHT: In this paper we propose a new algorithm, based on sparse convolutional coding, for detecting recurrent motifs of arbitrary structure up to a given length.
- 353, TITLE: Quantifying how much sensory information in a neural code is relevant for behavior
<https://papers.nips.cc/paper/6959-quantifying-how-much-sensory-information-in-a-neural-code-is-relevant-for-behavior>
AUTHORS: Giuseppe Pica, Eugenio Piasini, Houman Safaai, Caroline Runyan, Christopher Harvey, Mathew Diamond, Christoph Kayser, Tommaso Fellin, Stefano Panzeri
HIGHLIGHT: Here we develop a novel measure, termed the information-theoretic intersection information $\mathcal{M}(S;R;C)$, that quantifies how much of the sensory information carried by a neural response \mathcal{R} is used for behavior during perceptual discrimination tasks.
- 354, TITLE: Geometric Matrix Completion with Recurrent Multi-Graph Neural Networks
<https://papers.nips.cc/paper/6960-geometric-matrix-completion-with-recurrent-multi-graph-neural-networks>
AUTHORS: Federico Monti, Michael Bronstein, Xavier Bresson
HIGHLIGHT: We propose a novel approach to overcome these limitations by using geometric deep learning on graphs.
- 355, TITLE: Reducing Reparameterization Gradient Variance
<https://papers.nips.cc/paper/6961-reducing-reparameterization-gradient-variance>
AUTHORS: Andrew Miller, Nick Foti, Alexander D'Amour, Ryan P. Adams
HIGHLIGHT: We demonstrate our approach on a non-conjugate hierarchical model and a Bayesian neural net where our method attained orders of magnitude ($20 \times$) reduction in gradient variance resulting in faster and more stable optimization.
- 356, TITLE: Visual Reference Resolution using Attention Memory for Visual Dialog
<https://papers.nips.cc/paper/6962-visual-reference-resolution-using-attention-memory-for-visual-dialog>
AUTHORS: Paul Hongsuck Seo, Andreas Lehrmann, Bohyung Han, Leonid Sigal
HIGHLIGHT: We propose a novel attention mechanism that exploits visual attentions in the past to resolve the current reference in the visual dialog scenario.
- 357, TITLE: Joint distribution optimal transportation for domain adaptation

<https://papers.nips.cc/paper/6963-joint-distribution-optimal-transportation-for-domain-adaptation>

AUTHORS: Nicolas Courty, Rami Flamary, Amaury Habrard, Alain Rakotomamonjy
HIGHLIGHT: We propose a solution of this problem with optimal transport, that allows to recover an estimated target $\rho^* = (X, f(X))$ by optimizing simultaneously the optimal coupling and f .

358, TITLE: Multiresolution Kernel Approximation for Gaussian Process Regression

<https://papers.nips.cc/paper/6964-multiresolution-kernel-approximation-for-gaussian-process-regression>

AUTHORS: Yi Ding, Risi Kondor, Jonathan Eskreis-Winkler
HIGHLIGHT: In this paper we introduce Multiresolution Kernel Approximation (MKA), the first true broad bandwidth kernel approximation algorithm.

359, TITLE: Collapsed variational Bayes for Markov jump processes

<https://papers.nips.cc/paper/6965-collapsed-variational-bayes-for-markov-jump-processes>

AUTHORS: Boqian Zhang, Jiangwei Pan, Vinayak A. Rao
HIGHLIGHT: In this work, we propose a novel collapsed variational inference algorithm to address this issue.

360, TITLE: Universal consistency and minimax rates for online Mondrian Forests

<https://papers.nips.cc/paper/6966-universal-consistency-and-minimax-rates-for-online-mondrian-forests>

AUTHORS: Jaouad Mourtada, Stéphane Gaïffas, Erwan Scomet
HIGHLIGHT: We establish the consistency of an algorithm of Mondrian Forests \rightarrow cite {lakshminarayanan2014mondrianforests, lakshminarayanan2016mondrianuncertainty}, a randomized classification algorithm that can be implemented online.

361, TITLE: Welfare Guarantees from Data

<https://papers.nips.cc/paper/6967-welfare-guarantees-from-data>

AUTHORS: Darrell Hoy, Denis Nekipelov, Vasilis Syrgkanis
HIGHLIGHT: We propose a data-dependent analog of the price of anarchy that refines this worst-case assuming access to samples of strategic behavior.

362, TITLE: Diving into the shallows: a computational perspective on large-scale shallow learning

<https://papers.nips.cc/paper/6968-diving-into-the-shallows-a-computational-perspective-on-large-scale-shallow-learning>

AUTHORS: SIYUAN MA, Mikhail Belkin
HIGHLIGHT: In this paper we identify a basic limitation in gradient descent-based optimization methods when used in conjunctions with smooth kernels.

363, TITLE: End-to-end Differentiable Proving

<https://papers.nips.cc/paper/6969-end-to-end-differentiable-proving>

AUTHORS: Tim Rocktäschel, Sebastian Riedel
HIGHLIGHT: We introduce deep neural networks for end-to-end differentiable theorem proving that operate on dense vector representations of symbols.

364, TITLE: Influence Maximization with ϵ -Almost Submodular Threshold Functions

<https://papers.nips.cc/paper/6970-influence-maximization-with-epsilon-almost-submodular-threshold-functions>

AUTHORS: Qiang Li, Wei Chen, Institute of Computing Xiaoming Sun, Institute of Computing Jialin Zhang
HIGHLIGHT: In this paper, motivated by empirical evidences, we explore influence maximization in the non-submodular regime.

365, TITLE: InfoGAIL: Interpretable Imitation Learning from Visual Demonstrations

<https://papers.nips.cc/paper/6971-infogail-interpretable-imitation-learning-from-visual-demonstrations>

AUTHORS: Yunzhu Li, Jiaming Song, Stefano Ermon
HIGHLIGHT: In this paper, we propose a new algorithm that can infer the latent structure of expert demonstrations in an unsupervised way.

366, TITLE: Variational Laws of Visual Attention for Dynamic Scenes

<https://papers.nips.cc/paper/6972-variational-laws-of-visual-attention-for-dynamic-scenes>

AUTHORS: Dario Zanca, Marco Gori
HIGHLIGHT: This paper proposes a model of attentional scanpath that is based on the principle that there are foundational laws that drive the emergence of visual attention.

367, TITLE: Recursive Sampling for the Nystrom Method

<https://papers.nips.cc/paper/6973-recursive-sampling-for-the-nystrom-method>
AUTHORS: Cameron Musco, Christopher Musco
HIGHLIGHT: We give the first algorithm for kernel Nystrom approximation that runs in linear time in the number of training points and is provably accurate for all kernel matrices, without dependence on regularity or incoherence conditions.

368, TITLE: Interpolated Policy Gradient: Merging On-Policy and Off-Policy Gradient Estimation for Deep Reinforcement Learning
<https://papers.nips.cc/paper/6974-interpolated-policy-gradient-merging-on-policy-and-off-policy-gradient-estimation-for-deep-reinforcement-learning>
AUTHORS: Shixiang (Shane) Gu, Timothy Lillicrap, Richard E. Turner, Zoubin Ghahramani, Bernhard Schölkopf, Sergey Levine
HIGHLIGHT: This paper examines, both theoretically and empirically, approaches to merging on- and off-policy updates for deep reinforcement learning.

369, TITLE: Dynamic Routing Between Capsules
<https://papers.nips.cc/paper/6975-dynamic-routing-between-capsules>
AUTHORS: Sara Sabour, Nicholas Frosst, Geoffrey E. Hinton
HIGHLIGHT: We use the length of the activity vector to represent the probability that the entity exists and its orientation to represent the instantiation parameters.

370, TITLE: Incorporating Side Information by Adaptive Convolution
<https://papers.nips.cc/paper/6976-incorporating-side-information-by-adaptive-convolution>
AUTHORS: Di Kang, Debarun Dhar, Antoni Chan
HIGHLIGHT: In order to incorporate the available side information, we propose an adaptive convolutional neural network (ACNN), where the convolution filter weights adapt to the current scene context via the side information. Since existing crowd counting datasets do not contain ground-truth side information, we collect a new dataset with the ground-truth camera angle and height as the side information.

371, TITLE: Conic Scan-and-Cover algorithms for nonparametric topic modeling
<https://papers.nips.cc/paper/6977-conic-scan-and-cover-algorithms-for-nonparametric-topic-modeling>
AUTHORS: Mikhail Yurochkin, Aritra Guha, XuanLong Nguyen
HIGHLIGHT: We propose new algorithms for topic modeling when the number of topics is unknown.

372, TITLE: FALKON: An Optimal Large Scale Kernel Method
<https://papers.nips.cc/paper/6978-falkon-an-optimal-large-scale-kernel-method>
AUTHORS: Alessandro Rudi, Luigi Carratino, Lorenzo Rosasco
HIGHLIGHT: In this paper, we take a substantial step in scaling up kernel methods, proposing FALKON, a novel algorithm that allows to efficiently process millions of points.

373, TITLE: Structured Generative Adversarial Networks
<https://papers.nips.cc/paper/6979-structured-generative-adversarial-networks>
AUTHORS: Zhijie Deng, Hao Zhang, Xiaodan Liang, Luona Yang, Shizhen Xu, Jun Zhu, Eric P. Xing
HIGHLIGHT: We propose structured generative adversarial networks (SGANs) for semi-supervised conditional generative modeling.

374, TITLE: Conservative Contextual Linear Bandits
<https://papers.nips.cc/paper/6980-conservative-contextual-linear-bandits>
AUTHORS: Abbas Kazerouni, Mohammad Ghavamzadeh, Yasin Abbasi, Benjamin Van Roy
HIGHLIGHT: In this paper, we study the issue of safety in contextual linear bandits that have application in many different fields including personalized ad recommendation in online marketing.

375, TITLE: Variational Memory Addressing in Generative Models
<https://papers.nips.cc/paper/6981-variational-memory-addressing-in-generative-models>
AUTHORS: Jrg Bomschein, Andriy Mnih, Daniel Zoran, Danilo Jimenez Rezende
HIGHLIGHT: To illustrate the advantages of this approach we incorporate it into a variational autoencoder and apply the resulting model to the task of generative few-shot learning.

376, TITLE: On Tensor Train Rank Minimization : Statistical Efficiency and Scalable Algorithm
<https://papers.nips.cc/paper/6982-on-tensor-train-rank-minimization-statistical-efficiency-and-scalable-algorithm>
AUTHORS: Masaaki Imaizumi, Takanori Maehara, Kohei Hayashi
HIGHLIGHT: In this paper, we address the limitations.

- 377, TITLE: Scalable Levy Process Priors for Spectral Kernel Learning
<https://papers.nips.cc/paper/6983-scalable-levy-process-priors-for-spectral-kernel-learning>
AUTHORS: Phillip A. Jang, Andrew Loeb, Matthew Davidow, Andrew G. Wilson
HIGHLIGHT: We propose a distribution over kernels formed by modelling a spectral mixture density with a Levy process.
- 378, TITLE: Deep Hyperspherical Learning
<https://papers.nips.cc/paper/6984-deep-hyperspherical-learning>
AUTHORS: Weiyang Liu, Yan-Ming Zhang, Xingguo Li, Zhiding Yu, Bo Dai, Tuo Zhao, Le Song
HIGHLIGHT: In light of such challenges, we propose hyperspherical convolution (SphereConv), a novel learning framework that gives angular representations on hyperspheres.
- 379, TITLE: Learning Deep Structured Multi-Scale Features using Attention-Gated CRFs for Contour Prediction
<https://papers.nips.cc/paper/6985-learning-deep-structured-multi-scale-features-using-attention-gated-crf-for-contour-prediction>
AUTHORS: Dan Xu, Wanli Ouyang, Xavier Alameda-Pineda, Elisa Ricci, Xiaogang Wang, Nicu Sebe
HIGHLIGHT: This paper presents a novel approach for predicting contours which advances the state of the art in two fundamental aspects, i.e. multi-scale feature generation and fusion.
- 380, TITLE: On-the-fly Operation Batching in Dynamic Computation Graphs
<https://papers.nips.cc/paper/6986-on-the-fly-operation-batching-in-dynamic-computation-graphs>
AUTHORS: Graham Neubig, Yoav Goldberg, Chris Dyer
HIGHLIGHT: In this paper, we present an algorithm, and its implementation in the DyNet toolkit, for automatically batching operations.
- 381, TITLE: Nonlinear Acceleration of Stochastic Algorithms
<https://papers.nips.cc/paper/6987-nonlinear-acceleration-of-stochastic-algorithms>
AUTHORS: Damien Scieur, Francis Bach, Alexandre d'Aspremont
HIGHLIGHT: Here, we study extrapolation methods in a stochastic setting, where the iterates are produced by either a simple or an accelerated stochastic gradient algorithm.
- 382, TITLE: Optimized Pre-Processing for Discrimination Prevention
<https://papers.nips.cc/paper/6988-optimized-pre-processing-for-discrimination-prevention>
AUTHORS: Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, Kush R. Varshney
HIGHLIGHT: In this paper, we introduce a novel probabilistic formulation of data pre-processing for reducing discrimination.
- 383, TITLE: YASS: Yet Another Spike Sorter
<https://papers.nips.cc/paper/6989-yass-yet-another-spike-sorter>
AUTHORS: Jin Hyung Lee, David E. Carlson, Hooshmand Shokri Razaghi, Weichi Yao, Georges A. Goetz, Espen Hagen, Eleanor Batty, E.J. Chichilnisky, Gaute T. Einevoll, Liam Paninski
HIGHLIGHT: We present several new techniques that make dense MEA spike sorting more robust and scalable.
- 384, TITLE: Independence clustering (without a matrix)
<https://papers.nips.cc/paper/6990-independence-clustering-without-a-matrix>
AUTHORS: Daniil Ryabko
HIGHLIGHT: A consistent, computationally tractable algorithm for each of the settings is proposed, and a number of fascinating open directions for further research are outlined.
- 385, TITLE: Fast amortized inference of neural activity from calcium imaging data with variational autoencoders
<https://papers.nips.cc/paper/6991-fast-amortized-inference-of-neural-activity-from-calcium-imaging-data-with-variational-autoencoders>
AUTHORS: Artur Speiser, Jinyao Yan, Evan W. Archer, Lars Buesing, Srinivas C. Turaga, Jakob H. Macke
HIGHLIGHT: Here, we introduce a flexible algorithmic framework for fast, efficient and accurate extraction of neural spikes from imaging data.
- 386, TITLE: Adaptive Active Hypothesis Testing under Limited Information
<https://papers.nips.cc/paper/6992-adaptive-active-hypothesis-testing-under-limited-information>
AUTHORS: Fabio Cecchi, Nidhi Hegde
HIGHLIGHT: In this paper we consider a special case where the decision maker has limited knowledge about the distribution of observations for each action, in that only a binary value is observed.

- 387, TITLE: Streaming Weak Submodularity: Interpreting Neural Networks on the Fly
<https://papers.nips.cc/paper/6993-streaming-weak-submodularity-interpreting-neural-networks-on-the-fly>
AUTHORS: Ethan Elenberg, Alexandros G. Dimakis, Moran Feldman, Amin Karbasi
HIGHLIGHT: We cast interpretability of black-box classifiers as a combinatorial maximization problem and propose an efficient streaming algorithm to solve it subject to cardinality constraints.
- 388, TITLE: Successor Features for Transfer in Reinforcement Learning
<https://papers.nips.cc/paper/6994-successor-features-for-transfer-in-reinforcement-learning>
AUTHORS: Andre Barreto, Will Dabney, Remi Munos, Jonathan J. Hunt, Tom Schaul, Hado P. van Hasselt, David Silver
HIGHLIGHT: We propose a transfer framework for the scenario where the reward function changes between tasks but the environment's dynamics remain the same.
- 389, TITLE: Counterfactual Fairness
<https://papers.nips.cc/paper/6995-counterfactual-fairness>
AUTHORS: Matt J. Kusner, Joshua Loftus, Chris Russell, Ricardo Silva
HIGHLIGHT: In this paper, we develop a framework for modeling fairness using tools from causal inference.
- 390, TITLE: Prototypical Networks for Few-shot Learning
<https://papers.nips.cc/paper/6996-prototypical-networks-for-few-shot-learning>
AUTHORS: Jake Snell, Kevin Swersky, Richard Zemel
HIGHLIGHT: We propose Prototypical Networks for the problem of few-shot classification, where a classifier must generalize to new classes not seen in the training set, given only a small number of examples of each new class.
- 391, TITLE: Triple Generative Adversarial Nets
<https://papers.nips.cc/paper/6997-triple-generative-adversarial-nets>
AUTHORS: Chongxuan LI, Taufik Xu, Jun Zhu, Bo Zhang
HIGHLIGHT: To address the problems, we present triple generative adversarial net (Triple-GAN), which consists of three players---a generator, a discriminator and a classifier.
- 392, TITLE: Efficient Sublinear-Regret Algorithms for Online Sparse Linear Regression with Limited Observation
<https://papers.nips.cc/paper/6998-efficient-sublinear-regret-algorithms-for-online-sparse-linear-regression-with-limited-observation>
AUTHORS: Shinji Ito, Daisuke Hatano, Hanna Sumita, Akihiro Yabe, Takuro Fukunaga, Naonori Kakimura, Ken-Ichi Kawarabayashi
HIGHLIGHT: In this paper, we introduce mild assumptions to solve the problem.
- 393, TITLE: Mapping distinct timescales of functional interactions among brain networks
<https://papers.nips.cc/paper/6999-mapping-distinct-timescales-of-functional-interactions-among-brain-networks>
AUTHORS: Mali Sundaresan, Arshed Nabeel, Devarajan Sridharan
HIGHLIGHT: Here, we apply instantaneous and lag-based measures of conditional linear dependence, based on Granger-Geweke causality (GC), to infer network connections at distinct timescales from functional magnetic resonance imaging (fMRI) data.
- 394, TITLE: Multi-Armed Bandits with Metric Movement Costs
<https://papers.nips.cc/paper/7000-multi-armed-bandits-with-metric-movement-costs>
AUTHORS: Tomer Koren, Roi Livni, Yishay Mansour
HIGHLIGHT: Our main contribution gives a tight characterization of the expected minimax regret in this setting, in terms of a complexity measure C of the underlying metric which depends on its covering numbers.
- 395, TITLE: Learning A Structured Optimal Bipartite Graph for Co-Clustering
<https://papers.nips.cc/paper/7001-learning-a-structured-optimal-bipartite-graph-for-co-clustering>
AUTHORS: Feiping Nie, Xiaoqian Wang, Cheng Deng, Heng Huang
HIGHLIGHT: In this paper, we propose a novel co-clustering method to learn a bipartite graph with exactly k connected components, where k is the number of clusters.
- 396, TITLE: Learning Low-Dimensional Metrics
<https://papers.nips.cc/paper/7002-learning-low-dimensional-metrics>
AUTHORS: Blake Mason, Lalit Jain, Robert Nowak
HIGHLIGHT: This paper investigates the theoretical foundations of metric learning, focused on three key questions that are not fully addressed in prior work: 1) we consider learning general low-dimensional (low-rank) metrics as well as sparse metrics; 2) we develop upper and lower (minimax) bounds on the generalization error; 3) we quantify the sample complexity of metric learning in

terms of the dimension of the feature space and the dimension/rank of the underlying metric; 4) we also bound the accuracy of the learned metric relative to the underlying true generative metric.

- 397, TITLE: The Marginal Value of Adaptive Gradient Methods in Machine Learning
<https://papers.nips.cc/paper/7003-the-marginal-value-of-adaptive-gradient-methods-in-machine-learning>
AUTHORS: Ashia C. Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, Benjamin Recht
HIGHLIGHT: We construct an illustrative binary classification problem where the data is linearly separable, GD and SGD achieve zero test error, and AdaGrad, Adam, and RMSProp attain test errors arbitrarily close to half.
- 398, TITLE: Aggressive Sampling for Multi-class to Binary Reduction with Applications to Text Classification
<https://papers.nips.cc/paper/7004-aggressive-sampling-for-multi-class-to-binary-reduction-with-applications-to-text-classification>
AUTHORS: Bikash Joshi, Massih R. Amini, Ioannis Partalas, Franck Iutzeler, Yury Maximov
HIGHLIGHT: We propose a double sampling strategy on top of a multi-class to binary reduction strategy, which transforms the original multi-class problem into a binary classification problem over pairs of examples.
- 399, TITLE: Deconvolutional Paragraph Representation Learning
<https://papers.nips.cc/paper/7005-deconvolutional-paragraph-representation-learning>
AUTHORS: Yizhe Zhang, Dinghan Shen, Guoyin Wang, Zhe Gan, Ricardo Henao, Lawrence Carin
HIGHLIGHT: We propose a sequence-to-sequence, purely convolutional and deconvolutional autoencoding framework that is free of the above issue, while also being computationally efficient.
- 400, TITLE: Random Permutation Online Isotonic Regression
<https://papers.nips.cc/paper/7006-random-permutation-online-isotonic-regression>
AUTHORS: Wojciech Kotlowski, Wouter M. Koolen, Alan Malek
HIGHLIGHT: Here, we instead develop the more practical random permutation model.
- 401, TITLE: A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning
<https://papers.nips.cc/paper/7007-a-unified-game-theoretic-approach-to-multiagent-reinforcement-learning>
AUTHORS: Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Perolat, David Silver, Thore Graepel
HIGHLIGHT: We introduce a new metric, joint-policy correlation, to quantify this effect.
- 402, TITLE: Inverse Filtering for Hidden Markov Models
<https://papers.nips.cc/paper/7008-inverse-filtering-for-hidden-markov-models>
AUTHORS: Robert Mattila, Cristian Rojas, Vikram Krishnamurthy, Bo Wahlberg
HIGHLIGHT: This paper considers a number of related inverse filtering problems for hidden Markov models (HMMs).
- 403, TITLE: Non-parametric Structured Output Networks
<https://papers.nips.cc/paper/7009-non-parametric-structured-output-networks>
AUTHORS: Andreas Lehrmann, Leonid Sigal
HIGHLIGHT: We propose non-parametric structured output networks (NSON), a modular approach that cleanly separates a non-parametric, structured posterior representation from a discriminative inference scheme but allows joint end-to-end training of both components.
- 404, TITLE: Learning Active Learning from Data
<https://papers.nips.cc/paper/7010-learning-active-learning-from-data>
AUTHORS: Ksenia Konyushkova, Raphael Sznitman, Pascal Fua
HIGHLIGHT: In this paper, we suggest a novel data-driven approach to active learning (AL).
- 405, TITLE: VAE Learning via Stein Variational Gradient Descent
<https://papers.nips.cc/paper/7011-vae-learning-via-stein-variational-gradient-descent>
AUTHORS: Yuchen Pu, Zhe Gan, Ricardo Henao, Chunyuan Li, Shaobo Han, Lawrence Carin
HIGHLIGHT: A new method for learning variational autoencoders (VAEs) is developed, based on Stein variational gradient descent.
- 406, TITLE: Reconstructing perceived faces from brain activations with deep adversarial neural decoding
<https://papers.nips.cc/paper/7012-reconstructing-perceived-faces-from-brain-activations-with-deep-adversarial-neural-decoding>
AUTHORS: Yagmur G?l?rk, Umur G?l?, Katja Seeliger, Sander Bosch, Rob van Lier, Marcel A. J. van Gerven
HIGHLIGHT: Here, we present a novel approach to solve the problem of reconstructing perceived stimuli from brain responses by combining probabilistic inference with deep learning.

- 407, TITLE: Efficient Use of Limited-Memory Accelerators for Linear Learning on Heterogeneous Systems
<https://papers.nips.cc/paper/7013-efficient-use-of-limited-memory-accelerators-for-linear-learning-on-heterogeneous-systems>
AUTHORS: Celestine D'ner, Thomas Parnell, Martin Jaggi
HIGHLIGHT: We propose a generic algorithmic building block to accelerate training of machine learning models on heterogeneous compute systems.
- 408, TITLE: Temporal Coherency based Criteria for Predicting Video Frames using Deep Multi-stage Generative Adversarial Networks
<https://papers.nips.cc/paper/7014-temporal-coherency-based-criteria-for-predicting-video-frames-using-deep-multi-stage-generative-adversarial-networks>
AUTHORS: Prateep Bhattacharjee, Sukhendu Das
HIGHLIGHT: In this paper, we deal with this problem using convolutional models within a multi-stage Generative Adversarial Networks (GAN) framework.
- 409, TITLE: Sobolev Training for Neural Networks
<https://papers.nips.cc/paper/7015-sobolev-training-for-neural-networks>
AUTHORS: Wojciech M. Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Swirszcz, Razvan Pascanu
HIGHLIGHT: At the heart of deep learning we aim to use neural networks as function approximators - training them to produce outputs from inputs in emulation of a ground truth function or data creation process.
- 410, TITLE: Multi-Information Source Optimization
<https://papers.nips.cc/paper/7016-multi-information-source-optimization>
AUTHORS: Matthias Poloczek, Jialei Wang, Peter Frazier
HIGHLIGHT: We present a novel algorithm that outperforms the state of the art for this problem by using a Gaussian process covariance kernel better suited to MISO than those used by previous approaches, and an acquisition function based on a one-step optimality analysis supported by efficient parallelization.
- 411, TITLE: Deep Reinforcement Learning from Human Preferences
<https://papers.nips.cc/paper/7017-deep-reinforcement-learning-from-human-preferences>
AUTHORS: Paul F. Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, Dario Amodei
HIGHLIGHT: In this work, we explore goals defined in terms of (non-expert) human preferences between pairs of trajectory segments.
- 412, TITLE: On the Fine-Grained Complexity of Empirical Risk Minimization: Kernel Methods and Neural Networks
<https://papers.nips.cc/paper/7018-on-the-fine-grained-complexity-of-empirical-risk-minimization-kernel-methods-and-neural-networks>
AUTHORS: Arturs Backurs, Piotr Indyk, Ludwig Schmidt
HIGHLIGHT: We address this issue for multiple popular ERM problems including kernel SVMs, kernel ridge regression, and training the final layer of a neural network.
- 413, TITLE: Policy Gradient With Value Function Approximation For Collective Multiagent Planning
<https://papers.nips.cc/paper/7019-policy-gradient-with-value-function-approximation-for-collective-multiagent-planning>
AUTHORS: Duc Thien Nguyen, Akshat Kumar, Hoong Chuin Lau
HIGHLIGHT: Our main contribution is an actor-critic (AC) reinforcement learning method for optimizing CDec-POMDP policies.
- 414, TITLE: Adversarial Symmetric Variational Autoencoder
<https://papers.nips.cc/paper/7020-adversarial-symmetric-variational-autoencoder>
AUTHORS: Yuchen Pu, Weiyao Wang, Ricardo Henao, Liqun Chen, Zhe Gan, Chunyuan Li, Lawrence Carin
HIGHLIGHT: An extensive set of experiments is performed, in which we demonstrate state-of-the-art data reconstruction and generation on several image benchmarks datasets.
- 415, TITLE: Unified representation of tractography and diffusion-weighted MRI data using sparse multidimensional arrays
<https://papers.nips.cc/paper/7021-unified-representation-of-tractography-and-diffusion-weighted-mri-data-using-sparse-multidimensional-arrays>
AUTHORS: Cesar F. Caiafa, Olaf Sporns, Andrew Saykin, Franco Pestilli
HIGHLIGHT: In this paper, we introduce a method to encode dMRI signals and large connectomes, i.e., those that range from hundreds of thousands to millions of fascicles (bundles of neuronal axons), by using a sparse tensor decomposition.

- 416, TITLE: A Minimax Optimal Algorithm for Crowdsourcing
<https://papers.nips.cc/paper/7022-a-minimax-optimal-algorithm-for-crowdsourcing>
AUTHORS: Thomas Bonald, Richard Combes
HIGHLIGHT: We propose a novel lower bound on the minimax estimation error which applies to any estimation procedure.
- 417, TITLE: Estimating Accuracy from Unlabeled Data: A Probabilistic Logic Approach
<https://papers.nips.cc/paper/7023-estimating-accuracy-from-unlabeled-data-a-probabilistic-logic-approach>
AUTHORS: Emmanouil Platanios, Hoifung Poon, Tom M. Mitchell, Eric J. Horvitz
HIGHLIGHT: We propose an efficient method to estimate the accuracy of classifiers using only unlabeled data.
- 418, TITLE: A Decomposition of Forecast Error in Prediction Markets
<https://papers.nips.cc/paper/7024-a-decomposition-of-forecast-error-in-prediction-markets>
AUTHORS: Miro Dudik, Sebastien Lahaie, Ryan M. Rogers, Jennifer Wortman Vaughan
HIGHLIGHT: Our goal is to make explicit the tradeoffs between these error components, influenced by design decisions such as the functional form of the cost function and the amount of liquidity in the market.
- 419, TITLE: Safe Adaptive Importance Sampling
<https://papers.nips.cc/paper/7025-safe-adaptive-importance-sampling>
AUTHORS: Sebastian U. Stich, Anant Raj, Martin Jaggi
HIGHLIGHT: In this paper we propose an efficient approximation of gradient-based sampling, which is based on safe bounds on the gradient.
- 420, TITLE: Variational Walkback: Learning a Transition Operator as a Stochastic Recurrent Net
<https://papers.nips.cc/paper/7026-variational-walkback-learning-a-transition-operator-as-a-stochastic-recurrent-net>
AUTHORS: Anirudh Goyal ALIAS PARTH GOYAL, Nan Rosemary Ke, Surya Ganguli, Yoshua Bengio
HIGHLIGHT: We propose a novel method to $\{\text{it directly}\}$ learn a stochastic transition operator whose repeated application provides generated samples.
- 421, TITLE: Polynomial Codes: an Optimal Design for High-Dimensional Coded Matrix Multiplication
<https://papers.nips.cc/paper/7027-polynomial-codes-an-optimal-design-for-high-dimensional-coded-matrix-multiplication>
AUTHORS: Qian Yu, Mohammad Maddah-Ali, Salman Avestimehr
HIGHLIGHT: We propose a computation strategy that leverages ideas from coding theory to design intermediate computations at the worker nodes, in order to optimally deal with straggling workers.
- 422, TITLE: Unsupervised Learning of Disentangled Representations from Video
<https://papers.nips.cc/paper/7028-unsupervised-learning-of-disentangled-representations-from-video>
AUTHORS: Emily L. Denton, vighnesh Birodkar
HIGHLIGHT: We present a new model DRNET that learns disentangled image representations from video.
- 423, TITLE: Federated Multi-Task Learning
<https://papers.nips.cc/paper/7029-federated-multi-task-learning>
AUTHORS: Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, Ameet S. Talwalkar
HIGHLIGHT: In this work, we show that multi-task learning is naturally suited to handle the statistical challenges of this setting, and propose a novel systems-aware optimization method, MOCHA, that is robust to practical systems issues.
- 424, TITLE: Is Input Sparsity Time Possible for Kernel Low-Rank Approximation?
<https://papers.nips.cc/paper/7030-is-input-sparsity-time-possible-for-kernel-low-rank-approximation>
AUTHORS: Cameron Musco, David Woodruff
HIGHLIGHT: In this work we study the limits of computationally efficient low-rank kernel approximation.
- 425, TITLE: The Exp XORCIST: Nonparametric Graphical Models Via Conditional Exponential Densities
<https://papers.nips.cc/paper/7031-the-exp XORCIST-nonparametric-graphical-models-via-conditional-exponential-densities>
AUTHORS: Arun Suggala, Mladen Kolar, Pradeep K. Ravikumar
HIGHLIGHT: In this paper, we leverage recent developments to propose a class of non-parametric models which have very attractive computational and statistical properties.
- 426, TITLE: Improved Graph Laplacian via Geometric Self-Consistency
<https://papers.nips.cc/paper/7032-improved-graph-laplacian-via-geometric-self-consistency>
AUTHORS: Dominique Joncas, Marina Meila, James McQueen

HIGHLIGHT: We address the problem of setting the kernel bandwidth, epps, used by Manifold Learning algorithms to construct the graph Laplacian.

427, **TITLE:** Dual Path Networks

<https://papers.nips.cc/paper/7033-dual-path-networks>

AUTHORS: Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, Jiashi Feng

HIGHLIGHT: In this work, we present a simple, highly efficient and modularized Dual Path Network (DPN) for image classification which presents a new topology of connection paths internally.

428, **TITLE:** Faster and Non-ergodic $O(1/K)$ Stochastic Alternating Direction Method of Multipliers

<https://papers.nips.cc/paper/7034-faster-and-non-ergodic-o1k-stochastic-alternating-direction-method-of-multipliers>

AUTHORS: Cong Fang, Feng Cheng, Zhouchen Lin

HIGHLIGHT: In this paper, we propose a new stochastic ADMM which elaborately integrates Nesterov's extrapolation and VR techniques.

429, **TITLE:** A Probabilistic Framework for Nonlinearities in Stochastic Neural Networks

<https://papers.nips.cc/paper/7035-a-probabilistic-framework-for-nonlinearities-in-stochastic-neural-networks>

AUTHORS: Qinliang Su, xuejun Liao, Lawrence Carin

HIGHLIGHT: We present a probabilistic framework for nonlinearities, based on doubly truncated Gaussian distributions.

430, **TITLE:** Distral: Robust multitask reinforcement learning

<https://papers.nips.cc/paper/7036-distral-robust-multitask-reinforcement-learning>

AUTHORS: Yee Teh, Victor Bapst, Wojciech M. Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, Razvan Pascanu

HIGHLIGHT: We propose a new approach for joint training of multiple tasks, which we refer to as Distral (DISTill & TRansfer Learning).

431, **TITLE:** Online Learning of Optimal Bidding Strategy in Repeated Multi-Commodity Auctions

<https://papers.nips.cc/paper/7037-online-learning-of-optimal-bidding-strategy-in-repeated-multi-commodity-auctions>

AUTHORS: M. Sevi Baltaoglu, Lang Tong, Qing Zhao

HIGHLIGHT: As a bidding strategy, we propose a polynomial-time algorithm, inspired by the dynamic programming approach to the knapsack problem.

432, **TITLE:** Trimmed Density Ratio Estimation

<https://papers.nips.cc/paper/7038-trimmed-density-ratio-estimation>

AUTHORS: Song Liu, Akiko Takeda, Taiji Suzuki, Kenji Fukumizu

HIGHLIGHT: In this paper, we present a robust estimator which automatically identifies and trims outliers.

433, **TITLE:** Training recurrent networks to generate hypotheses about how the brain solves hard navigation problems

<https://papers.nips.cc/paper/7039-training-recurrent-networks-to-generate-hypotheses-about-how-the-brain-solves-hard-navigation-problems>

AUTHORS: Ingmar Kanitscheider, Ila Fiete

HIGHLIGHT: Our goal is to make sense of how the diverse phenomenology in the brain's spatial navigation circuits is related to their function.

434, **TITLE:** Visual Interaction Networks: Learning a Physics Simulator from Video

<https://papers.nips.cc/paper/7040-visual-interaction-networks-learning-a-physics-simulator-from-video>

AUTHORS: Nicholas Watters, Daniel Zoran, Theophane Weber, Peter Battaglia, Razvan Pascanu, Andrea Tacchetti

HIGHLIGHT: We introduce the Visual Interaction Network, a general-purpose model for learning the dynamics of a physical system from raw visual observations.

435, **TITLE:** Reconstruct & Crush Network

<https://papers.nips.cc/paper/7041-reconstruct-crush-network>

AUTHORS: Erinc Merdivan, Mohammad Reza Lohmani, Matthieu Geist

HIGHLIGHT: This article introduces an energy-based model that is adversarial regarding data: it minimizes the energy for a given data distribution (the positive samples) while maximizing the energy for another given data distribution (the negative or unlabeled samples).

436, **TITLE:** Streaming Robust Submodular Maximization: A Partitioned Thresholding Approach

<https://papers.nips.cc/paper/7042-streaming-robust-submodular-maximization-a-partitioned-thresholding-approach>

AUTHORS: Slobodan Mitrovic, Ilija Bogunovic, Ashkan Norouzi-Fard, Jakub M. Tarnawski, Volkan Cevher
HIGHLIGHT: We study the classical problem of maximizing a monotone submodular function subject to a cardinality constraint k , with two additional twists: (i) elements arrive in a streaming fashion, and (ii) m items from the algorithm's memory are removed after the stream is finished.

437, **TITLE:** Simple strategies for recovering inner products from coarsely quantized random projections
<https://papers.nips.cc/paper/7043-simple-strategies-for-recovering-inner-products-from-coarsely-quantized-random-projections>
AUTHORS: Ping Li, Martin Slawski
HIGHLIGHT: One implication is that in most scenarios of practical interest, there is no need for a sophisticated recovery approach like maximum likelihood estimation as considered in previous work on the subject.

438, **TITLE:** Discovering Potential Correlations via Hypercontractivity
<https://papers.nips.cc/paper/7044-discovering-potential-correlations-via-hypercontractivity>
AUTHORS: Hyeji Kim, Weihao Gao, Sreeram Kannan, Sewoong Oh, Pramod Viswanath
HIGHLIGHT: To bridge this gap, (i) we postulate a set of natural axioms that we expect a measure of potential correlation to satisfy; (ii) we show that the rate of information bottleneck, i.e., the hypercontractivity coefficient, satisfies all the proposed axioms; (iii) we provide a novel estimator to estimate the hypercontractivity coefficient from samples; and (iv) we provide numerical experiments demonstrating that this proposed estimator discovers potential correlations among various indicators of WHO datasets, is robust in discovering gene interactions from gene expression time series data, and is statistically more powerful than the estimators for other correlation measures in binary hypothesis testing of canonical examples of potential correlations.

439, **TITLE:** Doubly Stochastic Variational Inference for Deep Gaussian Processes
<https://papers.nips.cc/paper/7045-doubly-stochastic-variational-inference-for-deep-gaussian-processes>
AUTHORS: Hugh Salimbeni, Marc Deisenroth
HIGHLIGHT: We present a doubly stochastic variational inference algorithm, which does not force independence between layers.

440, **TITLE:** Ranking Data with Continuous Labels through Oriented Recursive Partitions
<https://papers.nips.cc/paper/7046-ranking-data-with-continuous-labels-through-oriented-recursive-partitions>
AUTHORS: Stéphane Clément, Mastane Achab
HIGHLIGHT: From the theoretical side, we describe the optimal elements of this problem and provide statistical guarantees for empirical Kendall τ maximization under appropriate conditions for the class of scoring function candidates.

441, **TITLE:** Scalable Model Selection for Belief Networks
<https://papers.nips.cc/paper/7047-scalable-model-selection-for-belief-networks>
AUTHORS: Zhao Song, Yusuke Muraoka, Ryohei Fujimaki, Lawrence Carin
HIGHLIGHT: We propose a scalable algorithm for model selection in sigmoid belief networks (SBNs), based on the factorized asymptotic Bayesian (FAB) framework.

442, **TITLE:** Targeting EEG/LFP Synchrony with Neural Nets
<https://papers.nips.cc/paper/7048-targeting-eeg-lfp-synchrony-with-neural-nets>
AUTHORS: Yitong Li, Michael Murias, Samantha Major, Geraldine Dawson, Kafui Dzirasa, Lawrence Carin, David E. Carlson
HIGHLIGHT: We consider the analysis of Electroencephalography (EEG) and Local Field Potential (LFP) datasets, which are "big" in terms of the size of recorded data but rarely have sufficient labels required to train complex models (e.g., conventional deep learning methods).

443, **TITLE:** Near-Optimal Edge Evaluation in Explicit Generalized Binomial Graphs
<https://papers.nips.cc/paper/7049-near-optimal-edge-evaluation-in-explicit-generalized-binomial-graphs>
AUTHORS: Sanjiban Choudhury, Shervin Javdani, Siddhartha Srinivasa, Sebastian Scherer
HIGHLIGHT: In this paper, we do so by drawing a novel equivalence between motion planning and the Bayesian active learning paradigm of decision region determination (DRD).

444, **TITLE:** Non-Stationary Spectral Kernels
<https://papers.nips.cc/paper/7050-non-stationary-spectral-kernels>
AUTHORS: Sami Remes, Markus Heinonen, Samuel Kaski
HIGHLIGHT: We propose non-stationary spectral kernels for Gaussian process regression by modelling the spectral density of a non-stationary kernel function as a mixture of input-dependent Gaussian process frequency density surfaces.

445, **TITLE:** Overcoming Catastrophic Forgetting by Incremental Moment Matching
<https://papers.nips.cc/paper/7051-overcoming-catastrophic-forgetting-by-incremental-moment-matching>

- AUTHORS: Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, Byoung-Tak Zhang
HIGHLIGHT: Here, we propose a method, i.e. incremental moment matching (IMM), to resolve this problem.
- 446, TITLE: Balancing information exposure in social networks
<https://papers.nips.cc/paper/7052-balancing-information-exposure-in-social-networks>
AUTHORS: Kiran Garimella, Aristides Gionis, Nikos Parotsidis, Nikolaj Tatti
HIGHLIGHT: In this paper we address the problem of balancing the information exposure} in a social network.
- 447, TITLE: SafetyNets: Verifiable Execution of Deep Neural Networks on an Untrusted Cloud
<https://papers.nips.cc/paper/7053-safetynets-verifiable-execution-of-deep-neural-networks-on-an-untrusted-cloud>
AUTHORS: Zahra Ghodsi, Tianyu Gu, Siddharth Garg
HIGHLIGHT: We propose SafetyNets, a framework that enables an untrusted server (the cloud) to provide a client with a short mathematical proof of the correctness of inference tasks that they perform on behalf of the client.
- 448, TITLE: Query Complexity of Clustering with Side Information
<https://papers.nips.cc/paper/7054-query-complexity-of-clustering-with-side-information>
AUTHORS: Arya Mazumdar, Barna Saha
HIGHLIGHT: In this paper, we provide a rigorous theoretical study of this basic problem of query complexity of interactive clustering, and give strong information theoretic lower bounds, as well as nearly matching upper bounds.
- 449, TITLE: QMDP-Net: Deep Learning for Planning under Partial Observability
<https://papers.nips.cc/paper/7055-qmdp-net-deep-learning-for-planning-under-partial-observability>
AUTHORS: Peter Karkus, David Hsu, Wee Sun Lee
HIGHLIGHT: This paper introduces the QMDP-net, a neural network architecture for planning under partial observability.
- 450, TITLE: Robust Optimization for Non-Convex Objectives
<https://papers.nips.cc/paper/7056-robust-optimization-for-non-convex-objectives>
AUTHORS: Robert S. Chen, Brendan Lucier, Yaron Singer, Vasilis Syrgkanis
HIGHLIGHT: We consider robust optimization problems, where the goal is to optimize in the worst case over a class of objective functions.
- 451, TITLE: Thy Friend is My Friend: Iterative Collaborative Filtering for Sparse Matrix Estimation
<https://papers.nips.cc/paper/7057-thy-friend-is-my-friend-iterative-collaborative-filtering-for-sparse-matrix-estimation>
AUTHORS: Christian Borgs, Jennifer Chayes, Christina E. Lee, Devavrat Shah
HIGHLIGHT: Inspired by classical collaborative filtering for recommendation systems, we propose a novel iterative, collaborative filtering-style algorithm for matrix estimation in this generic setting.
- 452, TITLE: Adaptive Classification for Prediction Under a Budget
<https://papers.nips.cc/paper/7058-adaptive-classification-for-prediction-under-a-budget>
AUTHORS: Feng Nan, Venkatesh Saligrama
HIGHLIGHT: We propose a novel adaptive approximation approach for test-time resource-constrained prediction motivated by Mobile, IoT, health, security and other applications, where constraints in the form of computation, communication, latency and feature acquisition costs arise.
We pose an empirical loss minimization problem with cost constraints to jointly train gating and prediction models.
- 453, TITLE: Convergence rates of a partition based Bayesian multivariate density estimation method
<https://papers.nips.cc/paper/7059-convergence-rates-of-a-partition-based-bayesian-multivariate-density-estimation-method>
AUTHORS: Linxi Liu, Dangna Li, Wing Hung Wong
HIGHLIGHT: We study a class of non-parametric density estimators under Bayesian settings.
- 454, TITLE: Affine-Invariant Online Optimization and the Low-rank Experts Problem
<https://papers.nips.cc/paper/7060-affine-invariant-online-optimization-and-the-low-rank-experts-problem>
AUTHORS: Tomer Koren, Roi Livni
HIGHLIGHT: We present a new affine-invariant optimization algorithm called Online Lazy Newton.
- 455, TITLE: Beyond Worst-case: A Probabilistic Analysis of Affine Policies in Dynamic Optimization
<https://papers.nips.cc/paper/7061-beyond-worst-case-a-probabilistic-analysis-of-affine-policies-in-dynamic-optimization>
AUTHORS: Omar El Housni, Vineet Goyal
HIGHLIGHT: In this paper, we aim to address this stark-contrast between the worst-case and the empirical performance of affine policies.

- 456, TITLE: A Unified Approach to Interpreting Model Predictions
<https://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions>
AUTHORS: Scott M. Lundberg, Su-In Lee
HIGHLIGHT: To address this problem, we present a unified framework for interpreting predictions, SHAP (SHapley Additive exPlanations).
- 457, TITLE: Stochastic Approximation for Canonical Correlation Analysis
<https://papers.nips.cc/paper/7063-stochastic-approximation-for-canonical-correlation-analysis>
AUTHORS: Raman Arora, Teodor Vanislavov Marinov, Poorya Mianjy, Nati Srebro
HIGHLIGHT: We propose novel first-order stochastic approximation algorithms for canonical correlation analysis (CCA).
- 458, TITLE: Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice
<https://papers.nips.cc/paper/7064-resurrecting-the-sigmoid-in-deep-learning-through-dynamical-isometry-theory-and-practice>
AUTHORS: Jeffrey Pennington, Samuel Schoenholz, Surya Ganguli
HIGHLIGHT: We address this question by employing powerful tools from free probability theory to analytically compute the singular value distribution of a deep network's input-output Jacobian.
- 459, TITLE: Sample and Computationally Efficient Learning Algorithms under S-Concave Distributions
<https://papers.nips.cc/paper/7065-sample-and-computationally-efficient-learning-algorithms-under-s-concave-distributions>
AUTHORS: Maria-Florina F. Balcan, Hongyang Zhang
HIGHLIGHT: In this work, we introduce new convex geometry tools to study the properties of S -concave distributions and use these properties to provide bounds on quantities of interest to learning including the probability of disagreement between two halfspaces, disagreement outside a band, and the disagreement coefficient.
- 460, TITLE: Scalable Variational Inference for Dynamical Systems
<https://papers.nips.cc/paper/7066-scalable-variational-inference-for-dynamical-systems>
AUTHORS: Nico S. Gorbach, Stefan Bauer, Joachim M. Buhmann
HIGHLIGHT: Using an existing gradient matching approach, we propose a scalable variational inference framework which can infer states and parameters simultaneously, offers computational speedups, improved accuracy and works well even under model misspecifications in a partially observable system.
- 461, TITLE: Context Selection for Embedding Models
<https://papers.nips.cc/paper/7067-context-selection-for-embedding-models>
AUTHORS: Liping Liu, Francisco Ruiz, Susan Athey, David Blei
HIGHLIGHT: In this paper, we show that conditioning on all the elements in the context is not optimal.
- 462, TITLE: Working hard to know your neighbor's margins: Local descriptor learning loss
<https://papers.nips.cc/paper/7068-working-hard-to-know-your-neighbors-margins-local-descriptor-learning-loss>
AUTHORS: Anastasiia Mishchuk, Dmytro Mishkin, Filip Radenovic, Jiri Matas
HIGHLIGHT: We introduce a loss for metric learning, which is inspired by the Lowe's matching criterion for SIFT.
- 463, TITLE: Accelerated Stochastic Greedy Coordinate Descent by Soft Thresholding Projection onto Simplex
<https://papers.nips.cc/paper/7069-accelerated-stochastic-greedy-coordinate-descent-by-soft-thresholding-projection-onto-simplex>
AUTHORS: Chaobing Song, Shaobo Cui, Yong Jiang, Shu-Tao Xia
HIGHLIGHT: In this paper we study the well-known greedy coordinate descent (GCD) algorithm to solve ℓ_1 -regularized problems and improve GCD by the two popular strategies: Nesterov's acceleration and stochastic optimization.
- 464, TITLE: Multi-Task Learning for Contextual Bandits
<https://papers.nips.cc/paper/7070-multi-task-learning-for-contextual-bandits>
AUTHORS: Aniket Anand Deshmukh, Urun Dogan, Clay Scott
HIGHLIGHT: In this work, we propose a multi-task learning framework for contextual bandit problems.
- 465, TITLE: Learning to Prune Deep Neural Networks via Layer-wise Optimal Brain Surgeon
<https://papers.nips.cc/paper/7071-learning-to-prune-deep-neural-networks-via-layer-wise-optimal-brain-surgeon>
AUTHORS: Xin Dong, Shangyu Chen, Sinno Pan
HIGHLIGHT: In this paper, we propose a new layer-wise pruning method for deep neural networks.
- 466, TITLE: Accelerated First-order Methods for Geodesically Convex Optimization on Riemannian Manifolds

<https://papers.nips.cc/paper/7072-accelerated-first-order-methods-for-geodesically-convex-optimization-on-riemannian-manifolds>
AUTHORS: Yuanyuan Liu, Fanhua Shang, James Cheng, Hong Cheng, Licheng Jiao
HIGHLIGHT: In this paper, we propose an accelerated first-order method for geodesically convex optimization, which is the generalization of the standard Nesterov's accelerated method from Euclidean space to nonlinear Riemannian space.

467, TITLE: Selective Classification for Deep Neural Networks
<https://papers.nips.cc/paper/7073-selective-classification-for-deep-neural-networks>
AUTHORS: Yonatan Geifman, Ran El-Yaniv
HIGHLIGHT: In this paper we propose a method to construct a selective classifier given a trained neural network.

468, TITLE: Minimax Estimation of Bandable Precision Matrices
<https://papers.nips.cc/paper/7074-minimax-estimation-of-bandable-precision-matrices>
AUTHORS: Addison Hu, Sahand Negahban
HIGHLIGHT: We establish minimax estimation bounds for estimating banded precision matrices under the spectral norm.

469, TITLE: Monte-Carlo Tree Search by Best Arm Identification
<https://papers.nips.cc/paper/7075-monte-carlo-tree-search-by-best-arm-identification>
AUTHORS: Emilie Kaufmann, Wouter M. Koolen
HIGHLIGHT: We study the game tree search problem, where the goal is to quickly identify the optimal move in a given game tree by sequentially sampling its stochastic payoffs.

470, TITLE: Group Additive Structure Identification for Kernel Nonparametric Regression
<https://papers.nips.cc/paper/7076-group-additive-structure-identification-for-kernel-nonparametric-regression>
AUTHORS: Chao Pan, Michael Zhu
HIGHLIGHT: The method utilizes a novel complexity measure we derive for group additive structures.

471, TITLE: Fast, Sample-Efficient Algorithms for Structured Phase Retrieval
<https://papers.nips.cc/paper/7077-fast-sample-efficient-algorithms-for-structured-phase-retrieval>
AUTHORS: Gauri Jagatap, Chinmay Hegde
HIGHLIGHT: We consider the problem of recovering a signal x in \mathbb{R}^n , from magnitude-only measurements, $y_i = |a_i^T x|$ for $i \in \{1, 2, \dots, m\}$.

472, TITLE: Hash Embeddings for Efficient Word Representations
<https://papers.nips.cc/paper/7078-hash-embeddings-for-efficient-word-representations>
AUTHORS: Dan Tito Svenstrup, Jonas Hansen, Ole Winther
HIGHLIGHT: We present hash embeddings, an efficient method for representing words in a continuous vector form.

473, TITLE: Online Learning for Multivariate Hawkes Processes
<https://papers.nips.cc/paper/7079-online-learning-for-multivariate-hawkes-processes>
AUTHORS: Yingxiang Yang, Jalal Etesami, Niao He, Negar Kiyavash
HIGHLIGHT: We develop a nonparametric and online learning algorithm that estimates the triggering functions of a multivariate Hawkes process (MHP).

474, TITLE: Maximum Margin Interval Trees
<https://papers.nips.cc/paper/7080-maximum-margin-interval-trees>
AUTHORS: Alexandre Drouin, Toby Hocking, Francois Laviolette
HIGHLIGHT: Whereas most existing algorithms for this task are linear models, in this paper we investigate learning nonlinear tree models.

475, TITLE: DropoutNet: Addressing Cold Start in Recommender Systems
<https://papers.nips.cc/paper/7081-dropoutnet-addressing-cold-start-in-recommender-systems>
AUTHORS: Maksims Volkovs, Guangwei Yu, Tomi Poutanen
HIGHLIGHT: Inspired by these results we propose a neural network based latent model called DropoutNet to address the cold start problem in recommender systems.

476, TITLE: A simple neural network module for relational reasoning
<https://papers.nips.cc/paper/7082-a-simple-neural-network-module-for-relational-reasoning>
AUTHORS: Adam Santoro, David Raposo, David G. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, Timothy Lillicrap

HIGHLIGHT: In this paper we describe how to use Relation Networks (RNs) as a simple plug-and-play module to solve problems that fundamentally hinge on relational reasoning.

477, **TITLE:** Q-LDA: Uncovering Latent Patterns in Text-based Sequential Decision Processes
<https://papers.nips.cc/paper/7083-q-lda-uncovering-latent-patterns-in-text-based-sequential-decision-processes>
AUTHORS: Jianshu Chen, Chong Wang, Lin Xiao, Ji He, Lihong Li, Li Deng
HIGHLIGHT: In this paper, we present a probabilistic model, Q-LDA, to uncover latent patterns in text-based sequential decision processes.

478, **TITLE:** Online Reinforcement Learning in Stochastic Games
<https://papers.nips.cc/paper/7084-online-reinforcement-learning-in-stochastic-games>
AUTHORS: Chen-Yu Wei, Yi-Te Hong, Chi-Jen Lu
HIGHLIGHT: We propose the \textsc{UCSG} algorithm that achieves a sublinear regret compared to the game value when competing with an arbitrary opponent.

479, **TITLE:** Position-based Multiple-play Bandit Problem with Unknown Position Bias
<https://papers.nips.cc/paper/7085-position-based-multiple-play-bandit-problem-with-unknown-position-bias>
AUTHORS: Junpei Komiyama, Junya Honda, Akiko Takeda
HIGHLIGHT: We propose the Permutation Minimum Empirical Divergence (PMED) algorithm and derive its asymptotically optimal regret bound.

480, **TITLE:** Active Exploration for Learning Symbolic Representations
<https://papers.nips.cc/paper/7086-active-exploration-for-learning-symbolic-representations>
AUTHORS: Garrett Andersen, George Konidaris
HIGHLIGHT: We introduce an online active exploration algorithm for data-efficiently learning an abstract symbolic model of an environment.

481, **TITLE:** Clone MCMC: Parallel High-Dimensional Gaussian Gibbs Sampling
<https://papers.nips.cc/paper/7087-clone-mcmc-parallel-high-dimensional-gaussian-gibbs-sampling>
AUTHORS: Andrei-Cristian Barbos, Francois Caron, Jean-François Giovannelli, Arnaud Doucet
HIGHLIGHT: We propose a generalized Gibbs sampler algorithm for obtaining samples approximately distributed from a high-dimensional Gaussian distribution.

482, **TITLE:** Fair Clustering Through Fairlets
<https://papers.nips.cc/paper/7088-fair-clustering-through-fairlets>
AUTHORS: Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, Sergei Vassilvitskii
HIGHLIGHT: En route we introduce the concept of fairlets, which are minimal sets that satisfy fair representation while approximately preserving the clustering objective.

483, **TITLE:** Polynomial time algorithms for dual volume sampling
<https://papers.nips.cc/paper/7089-polynomial-time-algorithms-for-dual-volume-sampling>
AUTHORS: Chengtao Li, Stefanie Jegelka, Suvrit Sra
HIGHLIGHT: This method was proposed by Avron and Boutsidis (2013), who showed it to be a promising method for column subset selection and its multiple applications.

484, **TITLE:** Hindsight Experience Replay
<https://papers.nips.cc/paper/7090-hindsight-experience-replay>
AUTHORS: Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, Wojciech Zaremba
HIGHLIGHT: We present a novel technique called Hindsight Experience Replay which allows sample-efficient learning from rewards which are sparse and binary and therefore avoid the need for complicated reward engineering.

485, **TITLE:** Stochastic and Adversarial Online Learning without Hyperparameters
<https://papers.nips.cc/paper/7091-stochastic-and-adversarial-online-learning-without-hyperparameters>
AUTHORS: Ashok Cutkosky, Kwabena A. Boahen
HIGHLIGHT: Here we introduce an online optimization algorithm that achieves $O(\log^4(T))$ regret in a wide class of stochastic settings while gracefully degrading to the optimal $O(\sqrt{T})$ regret in adversarial settings (up to logarithmic factors).

486, **TITLE:** Teaching Machines to Describe Images with Natural Language Feedback
<https://papers.nips.cc/paper/7092-teaching-machines-to-describe-images-with-natural-language-feedback>

- AUTHORS: huan ling, Sanja Fidler
HIGHLIGHT: In this paper, we bring a human in the loop, and enable a human teacher to give feedback to a learning agent in the form of natural language.
- 487, TITLE: Perturbative Black Box Variational Inference
<https://papers.nips.cc/paper/7093-perturbative-black-box-variational-inference>
AUTHORS: Robert Bamler, Cheng Zhang, Manfred Oppel, Stephan Mandt
HIGHLIGHT: In this paper, we view BBVI with generalized divergences as a form of estimating the marginal likelihood via biased importance sampling.
- 488, TITLE: GibbsNet: Iterative Adversarial Inference for Deep Graphical Models
<https://papers.nips.cc/paper/7094-gibbsnet-iterative-adversarial-inference-for-deep-graphical-models>
AUTHORS: Alex M. Lamb, Devon Hjelm, Yaroslav Ganin, Joseph Paul Cohen, Aaron C. Courville, Yoshua Bengio
HIGHLIGHT: We propose a novel approach to learning the joint distribution between the data and a latent code which uses an adversarially learned iterative procedure to gradually refine the joint distribution, $p(x, z)$, to better match with the data distribution on each step.
- 489, TITLE: PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space
<https://papers.nips.cc/paper/7095-pointnet-deep-hierarchical-feature-learning-on-point-sets-in-a-metric-space>
AUTHORS: Charles Ruizhongtai Qi, Li Yi, Hao Su, Leonidas J. Guibas
HIGHLIGHT: In this work, we introduce a hierarchical neural network that applies PointNet recursively on a nested partitioning of the input point set.
With further observation that point sets are usually sampled with varying densities, which results in greatly decreased performance for networks trained on uniform densities, we propose novel set learning layers to adaptively combine features from multiple scales.
- 490, TITLE: Regularizing Deep Neural Networks by Noise: Its Interpretation and Optimization
<https://papers.nips.cc/paper/7096-regularizing-deep-neural-networks-by-noise-its-interpretation-and-optimization>
AUTHORS: Hyeonwoo Noh, Tackgeun You, Jonghwan Mun, Bohyung Han
HIGHLIGHT: This paper addresses the above issues by 1) interpreting that the conventional training methods with regularization by noise injection optimize the lower bound of the true objective and 2) proposing a technique to achieve a tighter lower bound using multiple noise samples per training example in a stochastic gradient descent iteration.
- 491, TITLE: Learning Graph Representations with Embedding Propagation
<https://papers.nips.cc/paper/7097-learning-graph-representations-with-embedding-propagation>
AUTHORS: Alberto Garcia Duran, Mathias Niepert
HIGHLIGHT: We propose EP, Embedding Propagation, an unsupervised learning framework for graph-structured data.
- 492, TITLE: Efficient Modeling of Latent Information in Supervised Learning using Gaussian Processes
<https://papers.nips.cc/paper/7098-efficient-modeling-of-latent-information-in-supervised-learning-using-gaussian-processes>
AUTHORS: Zhenwen Dai, Mauricio Alvarez, Neil Lawrence
HIGHLIGHT: We present a new model called Latent Variable Multiple Output Gaussian Processes (LVMOGP) that allows to jointly model multiple conditions for regression and generalize to a new condition with a few data points at test time.
- 493, TITLE: A-NICE-MC: Adversarial Training for MCMC
<https://papers.nips.cc/paper/7099-a-nice-mc-adversarial-training-for-mcmc>
AUTHORS: Jiaming Song, Shengjia Zhao, Stefano Ermon
HIGHLIGHT: We propose A-NICE-MC, a novel method to train flexible parametric Markov chain kernels to produce samples with desired properties.
- 494, TITLE: Excess Risk Bounds for the Bayes Risk using Variational Inference in Latent Gaussian Models
<https://papers.nips.cc/paper/7100-excess-risk-bounds-for-the-bayes-risk-using-variational-inference-in-latent-gaussian-models>
AUTHORS: Rishit Sheth, Roni Khandon
HIGHLIGHT: To handle intractable inference, research in this area has developed new approximation methods that are fast and effective.
- 495, TITLE: Real-Time Bidding with Side Information
<https://papers.nips.cc/paper/7101-real-time-bidding-with-side-information>
AUTHORS: arthur flajolet, Patrick Jaillet
HIGHLIGHT: We consider the problem of repeated bidding in online advertising auctions when some side information (e.g. browser cookies) is available ahead of submitting a bid in the form of a d -dimensional vector.

- 496, TITLE: Saliency-based Sequential Image Attention with Multiset Prediction
<https://papers.nips.cc/paper/7102-saliency-based-sequential-image-attention-with-multiset-prediction>
AUTHORS: Sean Welleck, Jialin Mao, Kyunghyun Cho, Zheng Zhang
HIGHLIGHT: We propose a hierarchical visual architecture that operates on a saliency map and uses a novel attention mechanism to sequentially focus on salient regions and take additional glimpses within those regions.
- 497, TITLE: Variational Inference for Gaussian Process Models with Linear Complexity
<https://papers.nips.cc/paper/7103-variational-inference-for-gaussian-process-models-with-linear-complexity>
AUTHORS: Ching-An Cheng, Byron Boots
HIGHLIGHT: In this work, we propose a novel variational Gaussian process model that decouples the representation of mean and covariance functions in reproducing kernel Hilbert space.
- 498, TITLE: K-Medoids For K-Means Seeding
<https://papers.nips.cc/paper/7104-k-medoids-for-k-means-seeding>
AUTHORS: James Newling, François Fleuret
HIGHLIGHT: We introduce algorithmic improvements to CLARANS which improve its complexity and runtime, making it a viable initialization scheme for large datasets.
- 499, TITLE: Identifying Outlier Arms in Multi-Armed Bandit
<https://papers.nips.cc/paper/7105-identifying-outlier-arms-in-multi-armed-bandit>
AUTHORS: Honglei Zhuang, Chi Wang, Yifan Wang
HIGHLIGHT: We propose two sampling strategies with theoretical guarantee, and analyze their sampling efficiency.
- 500, TITLE: Online Learning with Transductive Regret
<https://papers.nips.cc/paper/7106-online-learning-with-transductive-regret>
AUTHORS: Mehryar Mohri, Scott Yang
HIGHLIGHT: We present a general and efficient online learning algorithm for minimizing transductive regret.
- 501, TITLE: Riemannian approach to batch normalization
<https://papers.nips.cc/paper/7107-riemannian-approach-to-batch-normalization>
AUTHORS: Minhyung Cho, Jaehyung Lee
HIGHLIGHT: Riemannian approach to batch normalization
- 502, TITLE: Self-supervised Learning of Motion Capture
<https://papers.nips.cc/paper/7108-self-supervised-learning-of-motion-capture>
AUTHORS: Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, Katerina Fragkiadaki
HIGHLIGHT: In this work, we propose a learning based motion capture model for single camera input.
- 503, TITLE: Triangle Generative Adversarial Networks
<https://papers.nips.cc/paper/7109-triangle-generative-adversarial-networks>
AUTHORS: Zhe Gan, Liqun Chen, Weiyao Wang, Yuchen Pu, Yizhe Zhang, Hao Liu, Chunyuan Li, Lawrence Carin
HIGHLIGHT: Triangle Generative Adversarial Networks
- 504, TITLE: PRUNE: Preserving Proximity and Global Ranking for Network Embedding
<https://papers.nips.cc/paper/7110-prune-preserving-proximity-and-global-ranking-for-network-embedding>
AUTHORS: Yi-An Lai, Chin-Chi Hsu, Wen Hao Chen, Mi-Yen Yeh, Shou-De Lin
HIGHLIGHT: A multi-task Siamese neural network structure is formulated to connect embedding vectors and our objective to preserve the global node ranking and local proximity of nodes.
- 505, TITLE: Bayesian Optimization with Gradients
<https://papers.nips.cc/paper/7111-bayesian-optimization-with-gradients>
AUTHORS: Jian Wu, Matthias Poloczek, Andrew G. Wilson, Peter Frazier
HIGHLIGHT: In this paper we show how Bayesian optimization can exploit derivative information to find good solutions with fewer objective function evaluations.
- 506, TITLE: Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation
<https://papers.nips.cc/paper/7112-scalable-trust-region-method-for-deep-reinforcement-learning-using-kronecker-factored-approximation>
AUTHORS: Yuhuai Wu, Elman Mansimov, Roger B. Grosse, Shun Liao, Jimmy Ba

HIGHLIGHT: In this work, we propose to apply trust region optimization to deep reinforcement learning using a recently proposed Kronecker-factored approximation to the curvature.

507, **TITLE:** Renyi Differential Privacy Mechanisms for Posterior Sampling
<https://papers.nips.cc/paper/7113-renyi-differential-privacy-mechanisms-for-posterior-sampling>
AUTHORS: Joseph Geumlek, Shuang Song, Kamalika Chaudhuri
HIGHLIGHT: We propose novel RDP mechanisms as well as offering a new RDP analysis for an existing method in order to add value to the RDP framework.

508, **TITLE:** Online Learning with a Hint
<https://papers.nips.cc/paper/7114-online-learning-with-a-hint>
AUTHORS: Ofer Dekel, arthur flajolet, Nika Haghtalab, Patrick Jaillet
HIGHLIGHT: We study a variant of online linear optimization where the player receives a hint about the loss function at the beginning of each round.

509, **TITLE:** Identification of Gaussian Process State Space Models
<https://papers.nips.cc/paper/7115-identification-of-gaussian-process-state-space-models>
AUTHORS: Stefanos Eleftheriadis, Tom Nicholson, Marc Deisenroth, James Hensman
HIGHLIGHT: We provide a practical algorithm for efficiently computing a lower bound on the marginal likelihood using the reparameterisation trick.

510, **TITLE:** Robust Imitation of Diverse Behaviors
<https://papers.nips.cc/paper/7116-robust-imitation-of-diverse-behaviors>
AUTHORS: Ziyu Wang, Josh S. Merel, Scott E. Reed, Nando de Freitas, Gregory Wayne, Nicolas Heess
HIGHLIGHT: In this paper, we show how to combine the favourable aspects of these two approaches.

511, **TITLE:** Can Decentralized Algorithms Outperform Centralized Algorithms? A Case Study for Decentralized Parallel Stochastic Gradient Descent
<https://papers.nips.cc/paper/7117-can-decentralized-algorithms-outperform-centralized-algorithms-a-case-study-for-decentralized-parallel-stochastic-gradient-descent>
AUTHORS: Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, Ji Liu
HIGHLIGHT: In this paper, we study a D-PSGD algorithm and provide the first theoretical analysis that indicates a regime in which decentralized algorithms might outperform centralized algorithms for distributed stochastic gradient descent.

512, **TITLE:** Local Aggregative Games
<https://papers.nips.cc/paper/7118-local-aggregative-games>
AUTHORS: Vikas Garg, Tommi Jaakkola
HIGHLIGHT: We focus on learning local aggregative games, where the payoff of each player is a function of its own action and the aggregate behavior of its neighbors in a connected digraph.

513, **TITLE:** A Sample Complexity Measure with Applications to Learning Optimal Auctions
<https://papers.nips.cc/paper/7119-a-sample-complexity-measure-with-applications-to-learning-optimal-auctions>
AUTHORS: Vasilis Syrgkanis
HIGHLIGHT: We introduce a new sample complexity measure, which we refer to as split-sample growth rate.

514, **TITLE:** Thinking Fast and Slow with Deep Learning and Tree Search
<https://papers.nips.cc/paper/7120-thinking-fast-and-slow-with-deep-learning-and-tree-search>
AUTHORS: Thomas Anthony, Zheng Tian, David Barber
HIGHLIGHT: In this paper, we present Expert Iteration (ExIt), a novel reinforcement learning algorithm which decomposes the problem into separate planning and generalisation tasks.

515, **TITLE:** EEG-GRAPH: A Factor-Graph-Based Model for Capturing Spatial, Temporal, and Observational Relationships in Electroencephalograms
<https://papers.nips.cc/paper/7121-eeg-graph-a-factor-graph-based-model-for-capturing-spatial-temporal-and-observational-relationships-in-electroencephalograms>
AUTHORS: Yogatheesan Varatharajah, Min Jin Chong, Krishnakant Saboo, Brent Berry, Benjamin Brinkmann, Gregory Worrell, Ravishankar Iyer
HIGHLIGHT: This paper presents a probabilistic-graphical model that can be used to infer characteristics of instantaneous brain activity by jointly analyzing spatial and temporal dependencies observed in electroencephalograms (EEG).

- 516, TITLE: Improving the Expected Improvement Algorithm
<https://papers.nips.cc/paper/7122-improving-the-expected-improvement-algorithm>
AUTHORS: Chao Qin, Diego Klabjan, Daniel Russo
HIGHLIGHT: To overcome this shortcoming, we introduce a simple modification of the expected improvement algorithm.
- 517, TITLE: Hybrid Reward Architecture for Reinforcement Learning
<https://papers.nips.cc/paper/7123-hybrid-reward-architecture-for-reinforcement-learning>
AUTHORS: Harm Van Seijen, Mehdi Fatemi, Joshua Romoff, Romain Laroche, Tavian Barnes, Jeffrey Tsang
HIGHLIGHT: This paper contributes towards tackling such challenging domains, by proposing a new method, called Hybrid Reward Architecture (HRA).
- 518, TITLE: Approximate Supermodularity Bounds for Experimental Design
<https://papers.nips.cc/paper/7124-approximate-supermodularity-bounds-for-experimental-design>
AUTHORS: Luiz Chamon, Alejandro Ribeiro
HIGHLIGHT: This work provides performance guarantees for the greedy solution of experimental design problems.
- 519, TITLE: Maximizing Subset Accuracy with Recurrent Neural Networks in Multi-label Classification
<https://papers.nips.cc/paper/7125-maximizing-subset-accuracy-with-recurrent-neural-networks-in-multi-label-classification>
AUTHORS: Jinseok Nam, Eneldo Loza Mencía, Hyunwoo J. Kim, Johannes Fürnkranz
HIGHLIGHT: In this paper, we replace classifier chains with recurrent neural networks, a sequence-to-sequence prediction algorithm which has recently been successfully applied to sequential prediction tasks in many domains.
- 520, TITLE: AdaGAN: Boosting Generative Models
<https://papers.nips.cc/paper/7126-adagan-boosting-generative-models>
AUTHORS: Ilya O. Tolstikhin, Sylvain Gelly, Olivier Bousquet, Carl-Johann SIMON-GABRIEL, Bernhard Schölkopf
HIGHLIGHT: We propose an iterative procedure, called AdaGAN, where at every step we add a new component into a mixture model by running a GAN algorithm on a re-weighted sample.
- 521, TITLE: Straggler Mitigation in Distributed Optimization Through Data Encoding
<https://papers.nips.cc/paper/7127-straggler-mitigation-in-distributed-optimization-through-data-encoding>
AUTHORS: Can Karakus, Yifan Sun, Suhas Diggavi, Wotao Yin
HIGHLIGHT: In this paper, we propose an alternate approach where we embed the redundancy directly in the data itself, and allow the computation to proceed completely oblivious to encoding.
- 522, TITLE: Multi-View Decision Processes: The Helper-AI Problem
<https://papers.nips.cc/paper/7128-multi-view-decision-processes-the-helper-ai-problem>
AUTHORS: Christos Dimitrakakis, David C. Parkes, Goran Radanovic, Paul Tylkin
HIGHLIGHT: We consider a two-player sequential game in which agents have the same reward function but may disagree on the transition probabilities of an underlying Markovian model of the world.
- 523, TITLE: A Greedy Approach for Budgeted Maximum Inner Product Search
<https://papers.nips.cc/paper/7129-a-greedy-approach-for-budgeted-maximum-inner-product-search>
AUTHORS: Hsiang-Fu Yu, Cho-Jui Hsieh, Qi Lei, Inderjit S. Dhillon
HIGHLIGHT: In this paper, we study the important problem of MIPS with a computational budget.
- 524, TITLE: SVD-Softmax: Fast Softmax Approximation on Large Vocabulary Neural Networks
<https://papers.nips.cc/paper/7130-svd-softmax-fast-softmax-approximation-on-large-vocabulary-neural-networks>
AUTHORS: Kyuhong Shim, Minjae Lee, Iksoo Choi, Yoonho Boo, Wonyong Sung
HIGHLIGHT: We propose a fast approximation method of a softmax function with a very large vocabulary using singular value decomposition (SVD).
- 525, TITLE: Plan, Attend, Generate: Planning for Sequence-to-Sequence Models
<https://papers.nips.cc/paper/7131-plan-attend-generate-planning-for-sequence-to-sequence-models>
AUTHORS: Caglar Gulcehre, Francis Dutil, Adam Trischler, Yoshua Bengio
HIGHLIGHT: We develop a model which can plan ahead in the future when it computes its alignments between input and output sequences, constructing a matrix of proposed future alignments and a commitment vector that governs whether to follow or recompute the plan.
- 526, TITLE: Task-based End-to-end Model Learning in Stochastic Optimization
<https://papers.nips.cc/paper/7132-task-based-end-to-end-model-learning-in-stochastic-optimization>

AUTHORS: Priya Donti, Brandon Amos, J. Zico Kolter
HIGHLIGHT: This paper proposes an end-to-end approach for learning probabilistic machine learning models in a manner that directly captures the ultimate task-based objective for which they will be used, within the context of stochastic programming.

527, TITLE: ALICE: Towards Understanding Adversarial Learning for Joint Distribution Matching
<https://papers.nips.cc/paper/7133-alice-towards-understanding-adversarial-learning-for-joint-distribution-matching>
AUTHORS: Chunyuan Li, Hao Liu, Changyou Chen, Yuchen Pu, Liqun Chen, Ricardo Henao, Lawrence Carin
HIGHLIGHT: Within a framework of conditional entropy, we propose both adversarial and non-adversarial approaches to learn desirable matched joint distributions for unsupervised and supervised tasks.

528, TITLE: Finite Sample Analysis of the GTD Policy Evaluation Algorithms in Markov Setting
<https://papers.nips.cc/paper/7134-finite-sample-analysis-of-the-gtd-policy-evaluation-algorithms-in-markov-setting>
AUTHORS: Yue Wang, Wei Chen, Yuting Liu, Zhi-Ming Ma, Tie-Yan Liu
HIGHLIGHT: In this paper, in the realistic Markov setting, we derive the finite sample bounds for the general convex-concave saddle point problems, and hence for the GTD algorithms.

529, TITLE: On the Complexity of Learning Neural Networks
<https://papers.nips.cc/paper/7135-on-the-complexity-of-learning-neural-networks>
AUTHORS: Le Song, Santosh Vempala, John Wilmes, Bo Xie
HIGHLIGHT: We demonstrate here a comprehensive lower bound ruling out this possibility: for a wide class of activation functions (including all currently used), and inputs drawn from any logconcave distribution, there is a family of one-hidden-layer functions whose output is a sum gate that are hard to learn in a precise sense: any statistical query algorithm (which includes all known variants of stochastic gradient descent with any loss function) needs an exponential number of queries even using tolerance inversely proportional to the input dimensionality.

530, TITLE: Hierarchical Implicit Models and Likelihood-Free Variational Inference
<https://papers.nips.cc/paper/7136-hierarchical-implicit-models-and-likelihood-free-variational-inference>
AUTHORS: Dustin Tran, Rajesh Ranganath, David Blei
HIGHLIGHT: HIMs combine the idea of implicit densities with hierarchical Bayesian modeling thereby defining models via simulators of data with rich hidden structure.

531, TITLE: Semi-supervised Learning with GANs: Manifold Invariance with Improved Inference
<https://papers.nips.cc/paper/7137-semi-supervised-learning-with-gans-manifold-invariance-with-improved-inference>
AUTHORS: Abhishek Kumar, Prasanna Sattigeri, Tom Fletcher
HIGHLIGHT: Motivated by the ability of the GANs generator to capture the data manifold well, we propose to estimate the tangent space to the data manifold using GANs and employ it to inject invariances into the classifier.

532, TITLE: Approximation and Convergence Properties of Generative Adversarial Learning
<https://papers.nips.cc/paper/7138-approximation-and-convergence-properties-of-generative-adversarial-learning>
AUTHORS: Shuang Liu, Olivier Bousquet, Kamalika Chaudhuri
HIGHLIGHT: In this paper, we address these questions in a broad and unified setting by defining a notion of adversarial divergences that includes a number of recently proposed objective functions.

533, TITLE: From Bayesian Sparsity to Gated Recurrent Nets
<https://papers.nips.cc/paper/7139-from-bayesian-sparsity-to-gated-recurrent-nets>
AUTHORS: Hao He, Bo Xin, Satoshi Ikehata, David Wipf
HIGHLIGHT: As part of this development, we examine the parallels between latent variable trajectories operating across multiple time-scales during optimization, and the activations within deep network structures designed to adaptively model such characteristic sequences.

534, TITLE: Min-Max Propagation
<https://papers.nips.cc/paper/7140-min-max-propagation>
AUTHORS: Christopher Srinivasa, Inmar Givoni, Siamak Ravanbakhsh, Brendan J. Frey
HIGHLIGHT: We study the application of min-max propagation, a variation of belief propagation, for approximate min-max inference in factor graphs.

535, TITLE: What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?
<https://papers.nips.cc/paper/7141-what-uncertainties-do-we-need-in-bayesian-deep-learning-for-computer-vision>
AUTHORS: Alex Kendall, Yarin Gal
HIGHLIGHT: For this we present a Bayesian deep learning framework combining input-dependent aleatoric uncertainty together with epistemic uncertainty.

- 536, TITLE: Gradient descent GAN optimization is locally stable
<https://papers.nips.cc/paper/7142-gradient-descent-gan-optimization-is-locally-stable>
AUTHORS: Vaishnavh Nagarajan, J. Zico Kolter
HIGHLIGHT: In this paper, we analyze the "gradient descent" form of GAN optimization, i.e., the natural setting where we simultaneously take small gradient steps in both generator and discriminator parameters.
- 537, TITLE: Toward Robustness against Label Noise in Training Deep Discriminative Neural Networks
<https://papers.nips.cc/paper/7143-toward-robustness-against-label-noise-in-training-deep-discriminative-neural-networks>
AUTHORS: Arash Vahdat
HIGHLIGHT: This paper proposes a novel framework for training deep convolutional neural networks from noisy labeled datasets that can be obtained cheaply.
- 538, TITLE: Dualing GANs
<https://papers.nips.cc/paper/7144-dualing-gans>
AUTHORS: Yujia Li, Alexander Schwing, Kuan-Chieh Wang, Richard Zemel
HIGHLIGHT: In this paper, we explore ways to tackle the instability problem by dualizing the discriminator.
- 539, TITLE: Deep Learning for Precipitation Nowcasting: A Benchmark and A New Model
<https://papers.nips.cc/paper/7145-deep-learning-for-precipitation-nowcasting-a-benchmark-and-a-new-model>
AUTHORS: Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, Wang-chun WOO
HIGHLIGHT: To address these problems, we propose both a new model and a benchmark for precipitation nowcasting. Besides, we provide a benchmark that includes a real-world large-scale dataset from the Hong Kong Observatory, a new training loss, and a comprehensive evaluation protocol to facilitate future research and gauge the state of the art.
- 540, TITLE: Do Deep Neural Networks Suffer from Crowding?
<https://papers.nips.cc/paper/7146-do-deep-neural-networks-suffer-from-crowding>
AUTHORS: Anna Volokitin, Gemma Roig, Tomaso A. Poggio
HIGHLIGHT: In this work, we study the effect of crowding in artificial Deep Neural Networks (DNNs) for object recognition.
- 541, TITLE: Learning from Complementary Labels
<https://papers.nips.cc/paper/7147-learning-from-complementary-labels>
AUTHORS: Takashi Ishida, Gang Niu, Weihua Hu, Masashi Sugiyama
HIGHLIGHT: In this paper, we show that an unbiased estimator to the classification risk can be obtained only from complementarily labeled data, if a loss function satisfies a particular symmetric condition.
- 542, TITLE: Online control of the false discovery rate with decaying memory
<https://papers.nips.cc/paper/7148-online-control-of-the-false-discovery-rate-with-decaying-memory>
AUTHORS: Aaditya Ramdas, Fanny Yang, Martin J. Wainwright, Michael I. Jordan
HIGHLIGHT: This paper improves the class of generalized alpha-investing algorithms (GAI) in four ways : (a) we show how to uniformly improve the power of the entire class of GAI procedures under independence by awarding more alpha-wealth for each rejection, giving a near win-win resolution to a dilemma raised by Javanmard and Montanari, (b) we demonstrate how to incorporate prior weights to indicate domain knowledge of which hypotheses are likely to be null or non-null, (c) we allow for differing penalties for false discoveries to indicate that some hypotheses may be more meaningful/important than others, (d) we define a new quantity called the "decaying memory false discovery rate, or memfdr " that may be more meaningful for applications with an explicit time component, using a discount factor to incrementally forget past decisions and alleviate some potential problems that we describe and name "piggybacking" and "alpha-death".
- 543, TITLE: Learning from uncertain curves: The 2-Wasserstein metric for Gaussian processes
<https://papers.nips.cc/paper/7149-learning-from-uncertain-curves-the-2-wasserstein-metric-for-gaussian-processes>
AUTHORS: Anton Mallasto, Aasa Feragen
HIGHLIGHT: We introduce a novel framework for statistical analysis of populations of non-degenerate Gaussian processes (GPs), which are natural representations of uncertain curves.
- 544, TITLE: Discriminative State Space Models
<https://papers.nips.cc/paper/7150-discriminative-state-space-models>
AUTHORS: Vitaly Kuznetsov, Mehryar Mohri
HIGHLIGHT: In this paper, we introduce and analyze Discriminative State-Space Models for forecasting non-stationary time series.

- 545, TITLE: On Fairness and Calibration
<https://papers.nips.cc/paper/7151-on-fairness-and-calibration>
AUTHORS: Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, Kilian Q. Weinberger
HIGHLIGHT: In this paper, we investigate the tension between minimizing error disparity across different population groups while maintaining calibrated probability estimates.
- 546, TITLE: Imagination-Augmented Agents for Deep Reinforcement Learning
<https://papers.nips.cc/paper/7152-imagination-augmented-agents-for-deep-reinforcement-learning>
AUTHORS: S?bastien Racani?re, Theophane Weber, David Reichert, Lars Buesing, Arthur Guez, Danilo Jimenez Rezende, Adri? Puigdom?nech Badia, Oriol Vinyals, Nicolas Heess, Yujia Li, Razvan Pascanu, Peter Battaglia, Demis Hassabis, David Silver, Daan Wierstra
HIGHLIGHT: We introduce Imagination-Augmented Agents (I2As), a novel architecture for deep reinforcement learning combining model-free and model-based aspects.
- 547, TITLE: Extracting low-dimensional dynamics from multiple large-scale neural population recordings by learning to predict correlations
<https://papers.nips.cc/paper/7153-extracting-low-dimensional-dynamics-from-multiple-large-scale-neural-population-recordings-by-learning-to-predict-correlations>
AUTHORS: Marcel Nonnenmacher, Srinivas C. Turaga, Jakob H. Macke
HIGHLIGHT: We propose an approach for extracting low-dimensional dynamics from multiple, sequential recordings.
- 548, TITLE: Unifying PAC and Regret: Uniform PAC Bounds for Episodic Reinforcement Learning
<https://papers.nips.cc/paper/7154-unifying-pac-and-regret-uniform-pac-bounds-for-episodic-reinforcement-learning>
AUTHORS: Christoph Dann, Tor Lattimore, Emma Brunskill
HIGHLIGHT: This paper introduces a new framework for theoretically measuring the performance of such algorithms called Uniform-PAC, which is a strengthening of the classical Probably Approximately Correct (PAC) framework.
- 549, TITLE: Gradients of Generative Models for Improved Discriminative Analysis of Tandem Mass Spectra
<https://papers.nips.cc/paper/7155-gradients-of-generative-models-for-improved-discriminative-analysis-of-tandem-mass-spectra>
AUTHORS: John T. Halloran, David M. Rocke
HIGHLIGHT: In this work, we leverage the log-likelihood gradients of generative models to improve the identification of such spectra.
- 550, TITLE: Asynchronous Parallel Coordinate Minimization for MAP Inference
<https://papers.nips.cc/paper/7156-asynchronous-parallel-coordinate-minimization-for-map-inference>
AUTHORS: Ofer Meshi, Alexander Schwing
HIGHLIGHT: In this work we propose to improve the efficiency of coordinate-minimization-based dual-decomposition solvers by running their updates asynchronously in parallel.
- 551, TITLE: Multiscale Quantization for Fast Similarity Search
<https://papers.nips.cc/paper/7157-multiscale-quantization-for-fast-similarity-search>
AUTHORS: Xiang Wu, Ruiqi Guo, Ananda Theertha Suresh, Sanjiv Kumar, Daniel N. Holtmann-Rice, David Simcha, Felix Yu
HIGHLIGHT: We propose a multiscale quantization approach for fast similarity search on large, high-dimensional datasets.
- 552, TITLE: Diverse and Accurate Image Description Using a Variational Auto-Encoder with an Additive Gaussian Encoding Space
<https://papers.nips.cc/paper/7158-diverse-and-accurate-image-description-using-a-variational-auto-encoder-with-an-additive-gaussian-encoding-space>
AUTHORS: Liwei Wang, Alexander Schwing, Svetlana Lazebnik
HIGHLIGHT: Instead, we propose two models that explicitly structure the latent space around K components corresponding to different types of image content, and combine components to create priors for images that contain multiple types of content simultaneously (e.g., several kinds of objects).
- 553, TITLE: Improved Training of Wasserstein GANs
<https://papers.nips.cc/paper/7159-improved-training-of-wasserstein-gans>
AUTHORS: Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, Aaron C. Courville
HIGHLIGHT: We propose an alternative to clipping weights: penalize the norm of gradient of the critic with respect to its input.
- 554, TITLE: Learning Populations of Parameters

<https://papers.nips.cc/paper/7160-learning-populations-of-parameters>

AUTHORS: Kevin Tian, Weihao Kong, Gregory Valiant

HIGHLIGHT: Consider the following estimation problem: there are n entities, each with an unknown parameter θ_i in $[0, 1]$, and we observe n independent random variables, X_1, \dots, X_n , with $X_i \sim \text{Binomial}(t, \theta_i)$.

555, TITLE: Clustering with Noisy Queries

<https://papers.nips.cc/paper/7161-clustering-with-noisy-queries>

AUTHORS: Arya Mazumdar, Barna Saha

HIGHLIGHT: In this paper, we provide a rigorous theoretical study of clustering with noisy queries.

556, TITLE: Higher-Order Total Variation Classes on Grids: Minimax Theory and Trend Filtering Methods

<https://papers.nips.cc/paper/7162-higher-order-total-variation-classes-on-grids-minimax-theory-and-trend-filtering-methods>

AUTHORS: Veeranjanyulu Sadhanala, Yu-Xiang Wang, James L. Sharpnack, Ryan J. Tibshirani

HIGHLIGHT: We consider the problem of estimating the values of a function over n nodes of a d -dimensional grid graph (having equal side lengths $n^{1/d}$) from noisy observations.

557, TITLE: Training Quantized Nets: A Deeper Understanding

<https://papers.nips.cc/paper/7163-training-quantized-nets-a-deeper-understanding>

AUTHORS: Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, Tom Goldstein

HIGHLIGHT: In this work, we investigate training methods for quantized neural networks from a theoretical viewpoint.

558, TITLE: Permutation-based Causal Inference Algorithms with Interventions

<https://papers.nips.cc/paper/7164-permutation-based-causal-inference-algorithms-with-interventions>

AUTHORS: Yuhao Wang, Liam Solus, Karren Yang, Caroline Uhler

HIGHLIGHT: In this paper, we present two algorithms of this type and prove that both are consistent under the faithfulness assumption.

559, TITLE: Time-dependent spatially varying graphical models, with application to brain fMRI data analysis

<https://papers.nips.cc/paper/7165-time-dependent-spatially-varying-graphical-models-with-application-to-brain-fmri-data-analysis>

AUTHORS: Kristjan Greenewald, Seyoung Park, Shuheng Zhou, Alexander Giessing

HIGHLIGHT: In this work, we present an additive model for space-time data that splits the data into a temporally correlated component and a spatially correlated component.

560, TITLE: Gradient Methods for Submodular Maximization

<https://papers.nips.cc/paper/7166-gradient-methods-for-submodular-maximization>

AUTHORS: Hamed Hassani, Mahdi Soltanolkotabi, Amin Karbasi

HIGHLIGHT: In this paper, we study the problem of maximizing continuous submodular functions that naturally arise in many learning applications such as those involving utility functions in active learning and sensing, matrix approximations and network inference.

561, TITLE: Smooth Primal-Dual Coordinate Descent Algorithms for Nonsmooth Convex Optimization

<https://papers.nips.cc/paper/7167-smooth-primal-dual-coordinate-descent-algorithms-for-nonsmooth-convex-optimization>

AUTHORS: Ahmet Alacaoglu, Quoc Tran Dinh, Olivier Fercoq, Volkan Cevher

HIGHLIGHT: We propose a new randomized coordinate descent method for a convex optimization template with broad applications.

562, TITLE: The Importance of Communities for Learning to Influence

<https://papers.nips.cc/paper/7168-the-importance-of-communities-for-learning-to-influence>

AUTHORS: Eric Balkanski, Nicole Immorlica, Yaron Singer

HIGHLIGHT: In this paper we describe a simple algorithm for maximizing influence from training data.

563, TITLE: Multiplicative Weights Update with Constant Step-Size in Congestion Games: Convergence, Limit Cycles and Chaos

<https://papers.nips.cc/paper/7169-multiplicative-weights-update-with-constant-step-size-in-congestion-games-convergence-limit-cycles-and-chaos>

AUTHORS: Gerasimos Palaiopoulos, Ioannis Panageas, Georgios Piliouras

HIGHLIGHT: We analyze MWU in congestion games where agents use ϵ as learning rates and prove convergence to exact Nash equilibria.

564, TITLE: Learning Neural Representations of Human Cognition across Many fMRI Studies

<https://papers.nips.cc/paper/7170-learning-neural-representations-of-human-cognition-across-many-fmri-studies>
AUTHORS: Arthur Mensch, Julien Mairal, Danilo Bzdok, Bertrand Thirion, Gael Varoquaux
HIGHLIGHT: We cast this challenge in a machine-learning approach to predict conditions from statistical brain maps across different studies.

565, TITLE: A KL-LUCB algorithm for Large-Scale Crowdsourcing
<https://papers.nips.cc/paper/7171-a-kl-lucb-algorithm-for-large-scale-crowdsourcing>
AUTHORS: Ervin Tanczos, Robert Nowak, Bob Mankoff
HIGHLIGHT: We develop an algorithm that is a fusion of lil-UCB and KL-LUCB, offering the best qualities of the two algorithms in one method.

566, TITLE: Collaborative Deep Learning in Fixed Topology Networks
<https://papers.nips.cc/paper/7172-collaborative-deep-learning-in-fixed-topology-networks>
AUTHORS: Zhanhong Jiang, Aditya Balu, Chinmay Hegde, Soumik Sarkar
HIGHLIGHT: In this context, this paper presents a new consensus-based distributed SGD (CDSGD) (and its momentum variant, CDMSGD) algorithm for collaborative deep learning over fixed topology networks that enables data parallelization as well as decentralized computation.

567, TITLE: Fast-Slow Recurrent Neural Networks
<https://papers.nips.cc/paper/7173-fast-slow-recurrent-neural-networks>
AUTHORS: Asier Mujika, Florian Meier, Angelika Steger
HIGHLIGHT: Here, we address this challenge by proposing a novel recurrent neural network (RNN) architecture, the Fast-Slow RNN (FS-RNN).

568, TITLE: Learning Disentangled Representations with Semi-Supervised Deep Generative Models
<https://papers.nips.cc/paper/7174-learning-disentangled-representations-with-semi-supervised-deep-generative-models>
AUTHORS: Siddharth Narayanaswamy, T. Brooks Paige, Jan-Willem van de Meent, Alban Desmaison, Noah Goodman, Pushmeet Kohli, Frank Wood, Philip Torr
HIGHLIGHT: We propose to learn such representations using model architectures that generalise from standard VAEs, employing a general graphical model structure in the encoder and decoder.

569, TITLE: Self-Supervised Intrinsic Image Decomposition
<https://papers.nips.cc/paper/7175-self-supervised-intrinsic-image-decomposition>
AUTHORS: Michael Janner, Jiajun Wu, Tejas D. Kulkarni, Ilker Yildirim, Josh Tenenbaum
HIGHLIGHT: In contrast to traditional fully supervised learning approaches, in this paper we propose learning intrinsic image decomposition by explaining the input image.

570, TITLE: Exploring Generalization in Deep Learning
<https://papers.nips.cc/paper/7176-exploring-generalization-in-deep-learning>
AUTHORS: Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, Nati Srebro
HIGHLIGHT: With a goal of understanding what drives generalization in deep networks, we consider several recently suggested explanations, including norm-based control, sharpness and robustness.

571, TITLE: A framework for Multi-A(rmed)/B(andit) Testing with Online FDR Control
<https://papers.nips.cc/paper/7177-a-framework-for-multi-armedbandit-testing-with-online-fdr-control>
AUTHORS: Fanny Yang, Aaditya Ramdas, Kevin G. Jamieson, Martin J. Wainwright
HIGHLIGHT: We propose an alternative framework to existing setups for controlling false alarms when multiple A/B tests are run over time.

572, TITLE: Fader Networks: Manipulating Images by Sliding Attributes
<https://papers.nips.cc/paper/7178-fader-networks-manipulating-images-by-sliding-attributes>
AUTHORS: Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic DENOYER, Marc'Aurelio Ranzato
HIGHLIGHT: This paper introduces a new encoder-decoder architecture that is trained to reconstruct images by disentangling the salient information of the image and the values of attributes directly in the latent space.

573, TITLE: Action Centered Contextual Bandits
<https://papers.nips.cc/paper/7179-action-centered-contextual-bandits>
AUTHORS: Kristjan Greenewald, Ambuj Tewari, Susan Murphy, Predag Klasnja
HIGHLIGHT: We provide an extension of the linear model for contextual bandits that has two parts: baseline reward and treatment effect.

- 574, TITLE: Estimating Mutual Information for Discrete-Continuous Mixtures
<https://papers.nips.cc/paper/7180-estimating-mutual-information-for-discrete-continuous-mixtures>
AUTHORS: Weihao Gao, Sreeram Kannan, Sewoong Oh, Pramod Viswanath
HIGHLIGHT: In this paper, we develop a novel estimator for estimating mutual information in discrete-continuous mixtures.
- 575, TITLE: Attention is All you Need
<https://papers.nips.cc/paper/7181-attention-is-all-you-need>
AUTHORS: Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin
HIGHLIGHT: We propose a novel, simple network architecture based solely on an attention mechanism, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train.
- 576, TITLE: Recurrent Ladder Networks
<https://papers.nips.cc/paper/7182-recurrent-ladder-networks>
AUTHORS: Isabeau Pr?mont-Schwarz, Alexander Ilin, Tele Hao, Antti Rasmus, Rinu Boney, Harri Valpola
HIGHLIGHT: We propose a recurrent extension of the Ladder networks whose structure is motivated by the inference required in hierarchical latent variable models.
- 577, TITLE: Parameter-Free Online Learning via Model Selection
<https://papers.nips.cc/paper/7183-parameter-free-online-learning-via-model-selection>
AUTHORS: Dylan J. Foster, Satyen Kale, Mehryar Mohri, Karthik Sridharan
HIGHLIGHT: We introduce an efficient algorithmic framework for model selection in online learning, also known as parameter-free online learning.
- 578, TITLE: Bregman Divergence for Stochastic Variance Reduction: Saddle-Point and Adversarial Prediction
<https://papers.nips.cc/paper/7184-bregman-divergence-for-stochastic-variance-reduction-saddle-point-and-adversarial-prediction>
AUTHORS: Zhan Shi, Xinhua Zhang, Yaoliang Yu
HIGHLIGHT: In this work we show that adversarial prediction under multivariate losses can be solved much faster than they used to be.
- 579, TITLE: Unbounded cache model for online language modeling with open vocabulary
<https://papers.nips.cc/paper/7185-unbounded-cache-model-for-online-language-modeling-with-open-vocabulary>
AUTHORS: Edouard Grave, Moustapha M. Cisse, Armand Joulin
HIGHLIGHT: In this paper, we propose an extension of continuous cache models, which can scale to larger contexts.
- 580, TITLE: Predictive State Recurrent Neural Networks
<https://papers.nips.cc/paper/7186-predictive-state-recurrent-neural-networks>
AUTHORS: Carlton Downey, Ahmed Hefny, Byron Boots, Geoffrey J. Gordon, Boyue Li
HIGHLIGHT: We present a new model, Predictive State Recurrent Neural Networks (PSRNNs), for filtering and prediction in dynamical systems.
- 581, TITLE: Early stopping for kernel boosting algorithms: A general analysis with localized complexities
<https://papers.nips.cc/paper/7187-early-stopping-for-kernel-boosting-algorithms-a-general-analysis-with-localized-complexities>
AUTHORS: Yuting Wei, Fanny Yang, Martin J. Wainwright
HIGHLIGHT: In this paper, for a relatively broad class of loss functions and boosting algorithms (including L^2 -boost, LogitBoost and AdaBoost, among others), we connect the performance of a stopped iterate to the localized Rademacher/Gaussian complexity of the associated function class.
- 582, TITLE: SVCCA: Singular Vector Canonical Correlation Analysis for Deep Learning Dynamics and Interpretability
<https://papers.nips.cc/paper/7188-svcca-singular-vector-canonical-correlation-analysis-for-deep-learning-dynamics-and-interpretability>
AUTHORS: Maithra Raghu, Justin Gilmer, Jason Yosinski, Jascha Sohl-Dickstein
HIGHLIGHT: We propose a new technique, Singular Vector Canonical Correlation Analysis (SVCCA), a tool for quickly comparing two representations in a way that is both invariant to affine transform (allowing comparison between different layers and networks) and fast to compute (allowing more comparisons to be calculated than with previous methods).
- 583, TITLE: Convolutional Phase Retrieval
<https://papers.nips.cc/paper/7189-convolutional-phase-retrieval>

- AUTHORS: Qing Qu, Yuqian Zhang, Yonina Eldar, John Wright
HIGHLIGHT: We study the convolutional phase retrieval problem, which asks us to recover an unknown signal \mathbf{x} of length n from m measurements consisting of the magnitude of its cyclic convolution with a known kernel \mathbf{a} of length m .
- 584, TITLE: Estimating High-dimensional Non-Gaussian Multiple Index Models via Stein's Lemma
<https://papers.nips.cc/paper/7190-estimating-high-dimensional-non-gaussian-multiple-index-models-via-steins-lemma>
AUTHORS: Zhuoran Yang, Krishnakumar Balasubramanian, Princeton Zhaoran Wang, Han Liu
HIGHLIGHT: To bypass the requirements of Gaussianity or elliptical symmetry of covariates in existing methods, we propose to leverage a second-order Stein's method with score function-based corrections.
- 585, TITLE: Gaussian Quadrature for Kernel Features
<https://papers.nips.cc/paper/7191-gaussian-quadrature-for-kernel-features>
AUTHORS: Tri Dao, Christopher M. De Sa, Christopher R?
HIGHLIGHT: The random Fourier features map is a technique commonly used to scale up kernel machines, but employing the randomized feature map means that $O(\epsilon^{-2})$ samples are required to achieve an approximation error of at most ϵ .
- 586, TITLE: Value Prediction Network
<https://papers.nips.cc/paper/7192-value-prediction-network>
AUTHORS: Junhyuk Oh, Satinder Singh, Honglak Lee
HIGHLIGHT: This paper proposes a novel deep reinforcement learning (RL) architecture, called Value Prediction Network (VPN), which integrates model-free and model-based RL methods into a single neural network.
- 587, TITLE: A Learning Error Analysis for Structured Prediction with Approximate Inference
<https://papers.nips.cc/paper/7193-a-learning-error-analysis-for-structured-prediction-with-approximate-inference>
AUTHORS: Yuanbin Wu, Man Lan, Shiliang Sun, Qi Zhang, Xuanjing Huang
HIGHLIGHT: In this work, we try to understand the differences between exact and approximate inference algorithms in structured prediction.
- 588, TITLE: Efficient Second-Order Online Kernel Learning with Adaptive Embedding
<https://papers.nips.cc/paper/7194-efficient-second-order-online-kernel-learning-with-adaptive-embedding>
AUTHORS: Daniele Calandriello, Alessandro Lazaric, Michal Valko
HIGHLIGHT: In this paper, we propose PROS-N-KONS, a method that combines Nystrom sketching to project the input point in a small, accurate embedded space, and performs efficient second-order updates in this space.
- 589, TITLE: Implicit Regularization in Matrix Factorization
<https://papers.nips.cc/paper/7195-implicit-regularization-in-matrix-factorization>
AUTHORS: Suriya Gunasekar, Blake E. Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, Nati Srebro
HIGHLIGHT: We study implicit regularization when optimizing an underdetermined quadratic objective over a matrix X with gradient descent on a factorization of X .
- 590, TITLE: Optimal Shrinkage of Singular Values Under Random Data Contamination
<https://papers.nips.cc/paper/7196-optimal-shrinkage-of-singular-values-under-random-data-contamination>
AUTHORS: Danny Barash, Matan Gavish
HIGHLIGHT: In this paper we show that common contamination models (including arbitrary combinations of uniform noise, missing values, outliers and corrupt entries) can be described efficiently using a single framework.
- 591, TITLE: Countering Feedback Delays in Multi-Agent Learning
<https://papers.nips.cc/paper/7197-countering-feedback-delays-in-multi-agent-learning>
AUTHORS: Zhengyuan Zhou, Panayotis Mertikopoulos, Nicholas Bambos, Peter W. Glynn, Claire Tomlin
HIGHLIGHT: We consider a model of game-theoretic learning based on online mirror descent (OMD) with asynchronous and delayed feedback information.
- 592, TITLE: Asynchronous Coordinate Descent under More Realistic Assumptions
<https://papers.nips.cc/paper/7198-asynchronous-coordinate-descent-under-more-realistic-assumptions>
AUTHORS: Tao Sun, Robert Hannah, Wotao Yin
HIGHLIGHT: In this paper, we argue that these assumptions either fail to hold or will imply less efficient implementations.
- 593, TITLE: Linear Convergence of a Frank-Wolfe Type Algorithm over Trace-Norm Balls
<https://papers.nips.cc/paper/7199-linear-convergence-of-a-frank-wolfe-type-algorithm-over-trace-norm-balls>

AUTHORS: Zeyuan Allen-Zhu, Elad Hazan, Wei Hu, Yuanzhi Li
HIGHLIGHT: We propose a rank-k variant of the classical Frank-Wolfe algorithm to solve convex optimization over a trace-norm ball.

594, TITLE: Hierarchical Clustering Beyond the Worst-Case
<https://papers.nips.cc/paper/7200-hierarchical-clustering-beyond-the-worst-case>
AUTHORS: Vincent Cohen-Addad, Varun Kanade, Frederik Mallmann-Trenn
HIGHLIGHT: In this paper, we consider a fairly general random graph model for hierarchical clustering, called the hierarchical stochastic blockmodel (HSBM), and show that in certain regimes the SVD approach of McSherry [5] combined with specific linkage methods results in a clustering that give an $O(1)$ -approximation to Dasgupta's cost function.

595, TITLE: Invariance and Stability of Deep Convolutional Representations
<https://papers.nips.cc/paper/7201-invariance-and-stability-of-deep-convolutional-representations>
AUTHORS: Alberto Bietti, Julien Mairal
HIGHLIGHT: In this paper, we study deep signal representations that are near-invariant to groups of transformations and stable to the action of diffeomorphisms without losing signal information.

596, TITLE: Statistical Cost Sharing
<https://papers.nips.cc/paper/7202-statistical-cost-sharing>
AUTHORS: Eric Balkanski, Umar Syed, Sergei Vassilvitskii
HIGHLIGHT: We study the cost sharing problem for cooperative games in situations where the cost function C is not available via oracle queries, but must instead be learned from samples drawn from a distribution, represented as tuples $(S, C(S))$, for different subsets S of players.

597, TITLE: The Expressive Power of Neural Networks: A View from the Width
<https://papers.nips.cc/paper/7203-the-expressive-power-of-neural-networks-a-view-from-the-width>
AUTHORS: Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, Liwei Wang
HIGHLIGHT: In this paper, we study how width affects the expressiveness of neural networks.

598, TITLE: Spectrally-normalized margin bounds for neural networks
<https://papers.nips.cc/paper/7204-spectrally-normalized-margin-bounds-for-neural-networks>
AUTHORS: Peter L. Bartlett, Dylan J. Foster, Matus J. Telgarsky
HIGHLIGHT: This paper presents a margin-based multiclass generalization bound for neural networks that scales with their margin-normalized "spectral complexity": their Lipschitz constant, meaning the product of the spectral norms of the weight matrices, times a certain correction factor.

599, TITLE: Robust and Efficient Transfer Learning with Hidden Parameter Markov Decision Processes
<https://papers.nips.cc/paper/7205-robust-and-efficient-transfer-learning-with-hidden-parameter-markov-decision-processes>
AUTHORS: Taylor W. Killian, Samuel Daulton, George Konidaris, Finale Doshi-Velez
HIGHLIGHT: We introduce a new formulation of the Hidden Parameter Markov Decision Process (HiP-MDP), a framework for modeling families of related tasks using low-dimensional latent embeddings.

600, TITLE: Population Matching Discrepancy and Applications in Deep Learning
<https://papers.nips.cc/paper/7206-population-matching-discrepancy-and-applications-in-deep-learning>
AUTHORS: Jianfei Chen, Chongxuan LI, Yizhong Ru, Jun Zhu
HIGHLIGHT: In this paper, we propose population matching discrepancy (PMD) for estimating the distribution distance based on samples, as well as an algorithm to learn the parameters of the distributions using PMD as an objective.

601, TITLE: Scalable Planning with Tensorflow for Hybrid Nonlinear Domains
<https://papers.nips.cc/paper/7207-scalable-planning-with-tensorflow-for-hybrid-nonlinear-domains>
AUTHORS: Ga Wu, Buser Say, Scott Sanner
HIGHLIGHT: Given recent deep learning results that demonstrate the ability to effectively optimize high-dimensional non-convex functions with gradient descent optimization on GPUs, we ask in this paper whether symbolic gradient optimization tools such as Tensorflow can be effective for planning in hybrid (mixed discrete and continuous) nonlinear domains with high dimensional state and action spaces?

602, TITLE: Boltzmann Exploration Done Right
<https://papers.nips.cc/paper/7208-boltzmann-exploration-done-right>
AUTHORS: Nicolò Cesa-Bianchi, Claudio Gentile, Gabor Lugosi, Gergely Neu
HIGHLIGHT: In this paper, we address several of these questions for the classic setup of stochastic multi-armed bandits.

- 603, TITLE: Learned in Translation: Contextualized Word Vectors
<https://papers.nips.cc/paper/7209-learned-in-translation-contextualized-word-vectors>
AUTHORS: Bryan McCann, James Bradbury, Caiming Xiong, Richard Socher
HIGHLIGHT: In this paper, we use a deep LSTM encoder from an attentional sequence-to-sequence model trained for machine translation (MT) to contextualize word vectors.
- 604, TITLE: Neural Discrete Representation Learning
<https://papers.nips.cc/paper/7210-neural-discrete-representation-learning>
AUTHORS: Aaron van den Oord, Oriol Vinyals, koray kavukcuoglu
HIGHLIGHT: In this paper, we propose a simple yet powerful generative model that learns such discrete representations.
- 605, TITLE: Generalizing GANs: A Turing Perspective
<https://papers.nips.cc/paper/7211-generalizing-gans-a-turing-perspective>
AUTHORS: Roderich Gross, Yue Gu, Wei Li, Melvin Gauci
HIGHLIGHT: In this paper we examine how these algorithms relate to the Turing test, and derive what - from a Turing perspective - can be considered their defining features.
- 606, TITLE: Scalable Log Determinants for Gaussian Process Kernel Learning
<https://papers.nips.cc/paper/7212-scalable-log-determinants-for-gaussian-process-kernel-learning>
AUTHORS: Kun Dong, David Eriksson, Hannes Nickisch, David Bindel, Andrew G. Wilson
HIGHLIGHT: We propose novel $O(n)$ approaches to estimating these quantities from only fast matrix vector multiplications (MVMs).
- 607, TITLE: Poincaré Embeddings for Learning Hierarchical Representations
<https://papers.nips.cc/paper/7213-poincare-embeddings-for-learning-hierarchical-representations>
AUTHORS: Maximillian Nickel, Douwe Kiela
HIGHLIGHT: In this work, we introduce a new approach for learning hierarchical representations of symbolic data by embedding them into hyperbolic space -- or more precisely into an n -dimensional Poincaré ball.
- 608, TITLE: Learning Combinatorial Optimization Algorithms over Graphs
<https://papers.nips.cc/paper/7214-learning-combinatorial-optimization-algorithms-over-graphs>
AUTHORS: Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, Le Song
HIGHLIGHT: In this paper, we propose a unique combination of reinforcement learning and graph embedding to address this challenge.
- 609, TITLE: Robust Conditional Probabilities
<https://papers.nips.cc/paper/7215-robust-conditional-probabilities>
AUTHORS: Yoav Wald, Amir Globerson
HIGHLIGHT: Here we propose a framework for reasoning about conditional probabilities without assuming anything about the underlying distributions, except knowledge of their second order marginals, which can be estimated from data.
- 610, TITLE: Learning with Bandit Feedback in Potential Games
<https://papers.nips.cc/paper/7216-learning-with-bandit-feedback-in-potential-games>
AUTHORS: Amélie Héliou, Johanne Cohen, Panayotis Mertikopoulos
HIGHLIGHT: In the bandit case, the same result holds for approximate Nash equilibria if we introduce a constant exploration factor that guarantees that action choice probabilities never become arbitrarily small.
- 611, TITLE: Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments
<https://papers.nips.cc/paper/7217-multi-agent-actor-critic-for-mixed-cooperative-competitive-environments>
AUTHORS: Ryan Lowe, YI WU, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, Igor Mordatch
HIGHLIGHT: We explore deep reinforcement learning methods for multi-agent domains.
- 612, TITLE: Communication-Efficient Distributed Learning of Discrete Distributions
<https://papers.nips.cc/paper/7218-communication-efficient-distributed-learning-of-discrete-distributions>
AUTHORS: Ilias Diakonikolas, Elena Grigorescu, Jerry Li, Abhiram Natarajan, Krzysztof Onak, Ludwig Schmidt
HIGHLIGHT: We initiate a systematic investigation of distribution learning (density estimation) when the data is distributed across multiple servers.
- 613, TITLE: Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles

- <https://papers.nips.cc/paper/7219-simple-and-scalable-predictive-uncertainty-estimation-using-deep-ensembles>
AUTHORS: Balaji Lakshminarayanan, Alexander Pritzel, Charles Blundell
HIGHLIGHT: We propose an alternative to Bayesian NNs that is simple to implement, readily parallelizable, requires very little hyperparameter tuning, and yields high quality predictive uncertainty estimates.
- 614, TITLE: When Worlds Collide: Integrating Different Counterfactual Assumptions in Fairness
<https://papers.nips.cc/paper/7220-when-worlds-collide-integrating-different-counterfactual-assumptions-in-fairness>
AUTHORS: Chris Russell, Matt J. Kusner, Joshua Loftus, Ricardo Silva
HIGHLIGHT: In this paper, we show how it is possible to make predictions that are approximately fair with respect to multiple possible causal models at once, thus mitigating the problem of exact causal specification.
- 615, TITLE: Matrix Norm Estimation from a Few Entries
<https://papers.nips.cc/paper/7221-matrix-norm-estimation-from-a-few-entries>
AUTHORS: Ashish Khetan, Sewoong Oh
HIGHLIGHT: We propose a framework of first estimating the Schatten k -norms of a matrix for several values of k , and using these as surrogates for estimating spectral properties of interest, such as the spectrum itself or the rank.
- 616, TITLE: Neural Networks for Efficient Bayesian Decoding of Natural Images from Retinal Neurons
<https://papers.nips.cc/paper/7222-neural-networks-for-efficient-bayesian-decoding-of-natural-images-from-retinal-neurons>
AUTHORS: Nikhil Parthasarathy, Eleanor Batty, William Falcon, Thomas Rutten, Mohit Rajpal, E.J. Chichilnisky, Liam Paninski
HIGHLIGHT: Here we develop a new approximate Bayesian method for decoding natural images from the spiking activity of populations of retinal ganglion cells (RGCs).
- 617, TITLE: Causal Effect Inference with Deep Latent-Variable Models
<https://papers.nips.cc/paper/7223-causal-effect-inference-with-deep-latent-variable-models>
AUTHORS: Christos Louizos, Uri Shalit, Joris M. Mooij, David Sontag, Richard Zemel, Max Welling
HIGHLIGHT: A carefully designed observational study attempts to measure all important confounders.
- 618, TITLE: Learning Identifiable Gaussian Bayesian Networks in Polynomial Time and Sample Complexity
<https://papers.nips.cc/paper/7224-learning-identifiable-gaussian-bayesian-networks-in-polynomial-time-and-sample-complexity>
AUTHORS: Asish Ghoshal, Jean Honorio
HIGHLIGHT: In this paper we propose a provably polynomial-time algorithm for learning sparse Gaussian Bayesian networks with equal noise variance --- a class of Bayesian networks for which the DAG structure can be uniquely identified from observational data --- under high-dimensional settings.
- 619, TITLE: Gradient Episodic Memory for Continual Learning
<https://papers.nips.cc/paper/7225-gradient-episodic-memory-for-continual-learning>
AUTHORS: David Lopez-Paz, Marc'Aurelio Ranzato
HIGHLIGHT: To better understand this issue, we study the problem of continual learning, where the model observes, once and one by one, examples concerning a sequence of tasks.
First, we propose a set of metrics to evaluate models learning over a continuum of data.
- 620, TITLE: Effective Parallelisation for Machine Learning
<https://papers.nips.cc/paper/7226-effective-parallelisation-for-machine-learning>
AUTHORS: Michael Kamp, Mario Boley, Olana Missura, Thomas G?rtner
HIGHLIGHT: We present a novel parallelisation scheme that simplifies the adaptation of learning algorithms to growing amounts of data as well as growing needs for accurate and confident predictions in critical applications.
- 621, TITLE: Semisupervised Clustering, AND-Queries and Locally Encodable Source Coding
<https://papers.nips.cc/paper/7227-semisupervised-clustering-and-queries-and-locally-encodable-source-coding>
AUTHORS: Arya Mazumdar, Soumyabrata Pal
HIGHLIGHT: In this paper, we show that a recently popular model of semisupervised clustering is equivalent to locally encodable source coding.
- 622, TITLE: Clustering Stable Instances of Euclidean k -means.
<https://papers.nips.cc/paper/7228-clustering-stable-instances-of-euclidean-k-means>
AUTHORS: Aravindan Vijayaraghavan, Abhratnu Dutta, Alex Wang
HIGHLIGHT: We design efficient algorithms that provably recover the optimal clustering for instances that are additive perturbation stable.

- 623, TITLE: Good Semi-supervised Learning That Requires a Bad GAN
<https://papers.nips.cc/paper/7229-good-semi-supervised-learning-that-requires-a-bad-gan>
AUTHORS: Zihang Dai, Zhilin Yang, Fan Yang, William W. Cohen, Ruslan R. Salakhutdinov
HIGHLIGHT: Theoretically we show that given the discriminator objective, good semi-supervised learning indeed requires a bad generator, and propose the definition of a preferred generator.
- 624, TITLE: On Blackbox Backpropagation and Jacobian Sensing
<https://papers.nips.cc/paper/7230-on-blackbox-backpropagation-and-jacobian-sensing>
AUTHORS: Krzysztof M. Choromanski, Vikas Sindhwani
HIGHLIGHT: From a small number of calls to a given "blackbox" on random input perturbations, we show how to efficiently recover its unknown Jacobian, or estimate the left action of its Jacobian on a given vector.
- 625, TITLE: Protein Interface Prediction using Graph Convolutional Networks
<https://papers.nips.cc/paper/7231-protein-interface-prediction-using-graph-convolutional-networks>
AUTHORS: Alex Fout, Jonathon Byrd, Basir Shariat, Asa Ben-Hur
HIGHLIGHT: We consider the prediction of interfaces between proteins, a challenging problem with important applications in drug discovery and design, and examine the performance of existing and newly proposed spatial graph convolution operators for this task.
- 626, TITLE: Solid Harmonic Wavelet Scattering: Predicting Quantum Molecular Energy from Invariant Descriptors of 3D Electronic Densities
<https://papers.nips.cc/paper/7232-solid-harmonic-wavelet-scattering-predicting-quantum-molecular-energy-from-invariant-descriptors-of-3d-electronic-densities>
AUTHORS: Michael Eickenberg, Georgios Exarchakis, Matthew Hirn, Stephane Mallat
HIGHLIGHT: We introduce a solid harmonic wavelet scattering representation, invariant to rigid motion and stable to deformations, for regression and classification of 2D and 3D signals.
- 627, TITLE: Towards Generalization and Simplicity in Continuous Control
<https://papers.nips.cc/paper/7233-towards-generalization-and-simplicity-in-continuous-control>
AUTHORS: Aravind Rajeswaran, Kendall Lowrey, Emanuel V. Todorov, Sham M. Kakade
HIGHLIGHT: The remarkable successes of deep learning in speech recognition and computer vision have motivated efforts to adapt similar techniques to other problem domains, including reinforcement learning (RL).
- 628, TITLE: Random Projection Filter Bank for Time Series Data
<https://papers.nips.cc/paper/7234-random-projection-filter-bank-for-time-series-data>
AUTHORS: Amir-massoud Farahmand, Sepideh Pourazarm, Daniel Nikovski
HIGHLIGHT: We propose Random Projection Filter Bank (RPFb) as a generic and simple approach to extract features from time series data.
- 629, TITLE: Filtering Variational Objectives
<https://papers.nips.cc/paper/7235-filtering-variational-objectives>
AUTHORS: Chris J. Maddison, John Lawson, George Tucker, Nicolas Heess, Mohammad Norouzi, Andriy Mnih, Arnaud Doucet, Yee Teh
HIGHLIGHT: We present results that relate the tightness of FIVO's bound to the variance of the particle filter's estimator by considering the generic case of bounds defined as log-transformed likelihood estimators.
- 630, TITLE: On Frank-Wolfe and Equilibrium Computation
<https://papers.nips.cc/paper/7236-on-frank-wolfe-and-equilibrium-computation>
AUTHORS: Jacob D. Abernethy, Jun-Kun Wang
HIGHLIGHT: We consider the Frank-Wolfe (FW) method for constrained convex optimization, and we show that this classical technique can be interpreted from a different perspective: FW emerges as the computation of an equilibrium (saddle point) of a special convex-concave zero sum game.
- 631, TITLE: Modulating early visual processing by language
<https://papers.nips.cc/paper/7237-modulating-early-visual-processing-by-language>
AUTHORS: Harm de Vries, Florian Strub, Jeremie Mary, Hugo Larochelle, Olivier Pietquin, Aaron C. Courville
HIGHLIGHT: In this paper, we deviate from this classic pipeline and propose to modulate the \emph{entire visual processing} by a linguistic input.
- 632, TITLE: Learning Mixture of Gaussians with Streaming Data

<https://papers.nips.cc/paper/7238-learning-mixture-of-gaussians-with-streaming-data>
AUTHORS: Aditi Raghunathan, Prateek Jain, Ravishankar Krishnawamy
HIGHLIGHT: In this paper, we study the problem of learning a mixture of Gaussians with streaming data: given a stream of S points in d dimensions generated by an unknown mixture of k spherical Gaussians, the goal is to estimate the model parameters using a single pass over the data stream.

633, TITLE: Practical Hash Functions for Similarity Estimation and Dimensionality Reduction
<https://papers.nips.cc/paper/7239-practical-hash-functions-for-similarity-estimation-and-dimensionality-reduction>
AUTHORS: Søren Dahlgaard, Mathias Knudsen, Mikkel Thorup
HIGHLIGHT: In this paper we focus on two prominent applications of hashing, namely similarity estimation with the one permutation hashing (OPH) scheme of Li et al. [NIPS'12] and feature hashing (FH) of Weinberger et al. [ICML'09], both of which have found numerous applications, i.e. in approximate near-neighbour search with LSH and large-scale classification with SVM.

634, TITLE: GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium
<https://papers.nips.cc/paper/7240-gans-trained-by-a-two-time-scale-update-rule-converge-to-a-local-nash-equilibrium>
AUTHORS: Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Sepp Hochreiter
HIGHLIGHT: We propose a two time-scale update rule (TTUR) for training GANs with stochastic gradient descent on arbitrary GAN loss functions.

635, TITLE: The Scaling Limit of High-Dimensional Online Independent Component Analysis
<https://papers.nips.cc/paper/7241-the-scaling-limit-of-high-dimensional-online-independent-component-analysis>
AUTHORS: Chuang Wang, Yue Lu
HIGHLIGHT: We analyze the dynamics of an online algorithm for independent component analysis in the high-dimensional scaling limit.

636, TITLE: Approximation Algorithms for ℓ_0 -Low Rank Approximation
https://papers.nips.cc/paper/7242-approximation-algorithms-for-ell_0-low-rank-approximation
AUTHORS: Karl Bringmann, Pavel Kolev, David Woodruff
HIGHLIGHT: We study the ℓ_0 -Low Rank Approximation Problem, where the goal is, given an $m \times n$ matrix A , to output a rank- k matrix A' for which $\|A' - A\|_0$ is minimized.

637, TITLE: The power of absolute discounting: all-dimensional distribution estimation
<https://papers.nips.cc/paper/7243-the-power-of-absolute-discounting-all-dimensional-distribution-estimation>
AUTHORS: Moein Falahatgar, Mesrob I. Ohannessian, Alon Orlitsky, Venkatesh Pichapati
HIGHLIGHT: In this paper, we rigorously explain the prowess of this estimator using less pessimistic notions.

638, TITLE: Few-Shot Adversarial Domain Adaptation
<https://papers.nips.cc/paper/7244-few-shot-adversarial-domain-adaptation>
AUTHORS: Saeid Motiian, Quinn Jones, Seyed Iranmanesh, Gianfranco Doretto
HIGHLIGHT: This work provides a framework for addressing the problem of supervised domain adaptation with deep models.

639, TITLE: Spectral Mixture Kernels for Multi-Output Gaussian Processes
<https://papers.nips.cc/paper/7245-spectral-mixture-kernels-for-multi-output-gaussian-processes>
AUTHORS: Gabriel Parra, Felipe Tobar
HIGHLIGHT: Spectral Mixture Kernels for Multi-Output Gaussian Processes

640, TITLE: Neural Expectation Maximization
<https://papers.nips.cc/paper/7246-neural-expectation-maximization>
AUTHORS: Klaus Greff, Sjoerd van Steenkiste, Jürgen Schmidhuber
HIGHLIGHT: In this paper, we explicitly formalize this problem as inference in a spatial mixture model where each component is parametrized by a neural network.

641, TITLE: Learning Linear Dynamical Systems via Spectral Filtering
<https://papers.nips.cc/paper/7247-learning-linear-dynamical-systems-via-spectral-filtering>
AUTHORS: Elad Hazan, Karan Singh, Cyril Zhang
HIGHLIGHT: We present an efficient and practical algorithm for the online prediction of discrete-time linear dynamical systems with a symmetric transition matrix.

642, TITLE: Z-Forcing: Training Stochastic Recurrent Networks
<https://papers.nips.cc/paper/7248-z-forcing-training-stochastic-recurrent-networks>

AUTHORS: Anirudh Goyal ALIAS PARTH GOYAL, Alessandro Sordoni, Marc-Alexandre C?t?, Nan Rosemary Ke, Yoshua Bengio
HIGHLIGHT: Many efforts have been devoted to training generative latent variable models with autoregressive decoders, such as recurrent neural networks (RNN).

643, TITLE: Learning Hierarchical Information Flow with Recurrent Neural Modules
<https://papers.nips.cc/paper/7249-learning-hierarchical-information-flow-with-recurrent-neural-modules>
AUTHORS: Danijar Hafner, Alexander Irpan, James Davidson, Nicolas Heess
HIGHLIGHT: We propose ThalNet, a deep learning model inspired by neocortical communication via the thalamus.

644, TITLE: Neural Variational Inference and Learning in Undirected Graphical Models
<https://papers.nips.cc/paper/7250-neural-variational-inference-and-learning-in-undirected-graphical-models>
AUTHORS: Volodymyr Kuleshov, Stefano Ermon
HIGHLIGHT: Here, we propose black-box learning and inference algorithms for undirected models that optimize a variational approximation to the log-likelihood of the model.

645, TITLE: Subspace Clustering via Tangent Cones
<https://papers.nips.cc/paper/7251-subspace-clustering-via-tangent-cones>
AUTHORS: Amin Jalali, Rebecca Willett
HIGHLIGHT: This paper proposes a new paradigm for subspace clustering that computes affinities based on the corresponding conic geometry.

646, TITLE: The Neural Hawkes Process: A Neurally Self-Modulating Multivariate Point Process
<https://papers.nips.cc/paper/7252-the-neural-hawkes-process-a-neurally-self-modulating-multivariate-point-process>
AUTHORS: Hongyuan Mei, Jason M. Eisner
HIGHLIGHT: Discovering such patterns can help us predict which type of event will happen next and when.

647, TITLE: Inverse Reward Design
<https://papers.nips.cc/paper/7253-inverse-reward-design>
AUTHORS: Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart J. Russell, Anca Dragan
HIGHLIGHT: We introduce approximate methods for solving IRD problems, and use their solution to plan risk-averse behavior in test MDPs.

648, TITLE: Structured Bayesian Pruning via Log-Normal Multiplicative Noise
<https://papers.nips.cc/paper/7254-structured-bayesian-pruning-via-log-normal-multiplicative-noise>
AUTHORS: Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, Dmitry P. Vetrov
HIGHLIGHT: In the paper, we propose a new Bayesian model that takes into account the computational structure of neural networks and provides structured sparsity, e.g. removes neurons and/or convolutional channels in CNNs.

649, TITLE: Attend and Predict: Understanding Gene Regulation by Selective Attention on Chromatin
<https://papers.nips.cc/paper/7255-attend-and-predict-understanding-gene-regulation-by-selective-attention-on-chromatin>
AUTHORS: Ritambhara Singh, Jack Lanchantin, Arshdeep Sekhon, Yanjun Qi
HIGHLIGHT: This paper presents an attention-based deep learning approach; AttentiveChrome, that uses a unified architecture to model and to interpret dependencies among chromatin factors for controlling gene regulation.

650, TITLE: Acceleration and Averaging in Stochastic Descent Dynamics
<https://papers.nips.cc/paper/7256-acceleration-and-averaging-in-stochastic-descent-dynamics>
AUTHORS: Walid Krichene, Peter L. Bartlett
HIGHLIGHT: Building on an averaging formulation of accelerated mirror descent, we propose a stochastic variant in which the gradient is contaminated by noise, and study the resulting stochastic differential equation.

651, TITLE: Kernel functions based on triplet comparisons
<https://papers.nips.cc/paper/7257-kernel-functions-based-on-triplet-comparisons>
AUTHORS: Matthias Kleindessner, Ulrike von Luxburg
HIGHLIGHT: Given only information in the form of similarity triplets "Object A is more similar to object B than to object C" about a data set, we propose two ways of defining a kernel function on the data set.

652, TITLE: An Error Detection and Correction Framework for Connectomics
<https://papers.nips.cc/paper/7258-an-error-detection-and-correction-framework-for-connectomics>
AUTHORS: Jonathan Zung, Ignacio Tartavull, Kisuk Lee, H. Sebastian Seung

HIGHLIGHT: We define and study error detection and correction tasks that are useful for 3D reconstruction of neurons from electron microscopic imagery, and for image segmentation more generally.

653, **TITLE:** Style Transfer from Non-Parallel Text by Cross-Alignment
<https://papers.nips.cc/paper/7259-style-transfer-from-non-parallel-text-by-cross-alignment>

AUTHORS: Tianxiao Shen, Tao Lei, Regina Barzilay, Tommi Jaakkola

HIGHLIGHT: We assume a shared latent content distribution across different text corpora, and propose a method that leverages refined alignment of latent representations to perform style transfer.

654, **TITLE:** Cross-Spectral Factor Analysis
<https://papers.nips.cc/paper/7260-cross-spectral-factor-analysis>

AUTHORS: Neil Gallagher, Kyle R. Ulrich, Austin Talbot, Kafui Dzirasa, Lawrence Carin, David E. Carlson

HIGHLIGHT: To facilitate understanding of network-level synchronization between brain regions, we introduce a novel model of multisite low-frequency neural recordings, such as local field potentials (LFPs) and electroencephalograms (EEGs).

655, **TITLE:** Stochastic Submodular Maximization: The Case of Coverage Functions
<https://papers.nips.cc/paper/7261-stochastic-submodular-maximization-the-case-of-coverage-functions>

AUTHORS: Mohammad Karimi, Mario Lucic, Hamed Hassani, Andreas Krause

HIGHLIGHT: We seek to unleash the power of stochastic continuous optimization, namely stochastic gradient descent and its variants, to such discrete problems.

We first introduce the problem of stochastic submodular optimization, where one needs to optimize a submodular objective which is given as an expectation.

656, **TITLE:** Affinity Clustering: Hierarchical Clustering at Scale
<https://papers.nips.cc/paper/7262-affinity-clustering-hierarchical-clustering-at-scale>

AUTHORS: Mohammadhossein Bateni, Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Raimondas Kiveris, Silvio Lattanzi, Vahab Mirrokni

HIGHLIGHT: In particular, we propose affinity, a novel hierarchical clustering based on Boruvka's MST algorithm.

657, **TITLE:** Unsupervised Transformation Learning via Convex Relaxations
<https://papers.nips.cc/paper/7263-unsupervised-transformation-learning-via-convex-relaxations>

AUTHORS: Tatsunori B. Hashimoto, Percy S. Liang, John C. Duchi

HIGHLIGHT: We propose an unsupervised approach to learn such transformations by attempting to reconstruct an image from a linear combination of transformations of its nearest neighbors.

658, **TITLE:** A Sharp Error Analysis for the Fused Lasso, with Application to Approximate Change-point Screening
<https://papers.nips.cc/paper/7264-a-sharp-error-analysis-for-the-fused-lasso-with-application-to-approximate-change-point-screening>

AUTHORS: Kevin Lin, James L. Sharpnack, Alessandro Rinaldo, Ryan J. Tibshirani

HIGHLIGHT: In the 1-dimensional multiple change-point detection problem, we derive a new fast error rate for the fused lasso estimator, under the assumption that the mean vector has a sparse number of change-points.

659, **TITLE:** Linear Time Computation of Moments in Sum-Product Networks
<https://papers.nips.cc/paper/7265-linear-time-computation-of-moments-in-sum-product-networks>

AUTHORS: Han Zhao, Geoffrey J. Gordon

HIGHLIGHT: We propose an optimal linear-time algorithm that works even when the SPN is a general directed acyclic graph (DAG), which significantly broadens the applicability of Bayesian online algorithms for SPNs.

660, **TITLE:** A Meta-Learning Perspective on Cold-Start Recommendations for Items
<https://papers.nips.cc/paper/7266-a-meta-learning-perspective-on-cold-start-recommendations-for-items>

AUTHORS: Manasi Vartak, Arvind Thiagarajan, Conrado Miranda, Jeshua Bratman, Hugo Larochelle

HIGHLIGHT: In this paper, we present a meta-learning strategy to address item cold-start when new items arrive continuously.

661, **TITLE:** Predicting Scene Parsing and Motion Dynamics in the Future
<https://papers.nips.cc/paper/7267-predicting-scene-parsing-and-motion-dynamics-in-the-future>

AUTHORS: Xiaojie Jin, Huaxin Xiao, Xiaohui Shen, Jimei Yang, Zhe Lin, Yunpeng Chen, Zequn Jie, Jiashi Feng, Shuicheng Yan

HIGHLIGHT: In this paper, we propose a novel model to predict the scene parsing and motion dynamics in unobserved future video frames simultaneously.

- 662, TITLE: Sticking the Landing: Simple, Lower-Variance Gradient Estimators for Variational Inference
<https://papers.nips.cc/paper/7268-sticking-the-landing-simple-lower-variance-gradient-estimators-for-variational-inference>
AUTHORS: Geoffrey Roeder, Yuhuai Wu, David K. Duvenaud
HIGHLIGHT: We propose a simple and general variant of the standard reparameterized gradient estimator for the variational evidence lower bound.
- 663, TITLE: Efficient Approximation Algorithms for Strings Kernel Based Sequence Classification
<https://papers.nips.cc/paper/7269-efficient-approximation-algorithms-for-strings-kernel-based-sequence-classification>
AUTHORS: Muhammad Farhan, Juvaria Tariq, Arif Zaman, Mudassir Shabbir, Imdad Ullah Khan
HIGHLIGHT: In this work, we develop novel techniques to efficiently and accurately estimate the pairwise similarity score, which enables us to use much larger values of k and m , and get higher predictive accuracy.
- 664, TITLE: Kernel Feature Selection via Conditional Covariance Minimization
<https://papers.nips.cc/paper/7270-kernel-feature-selection-via-conditional-covariance-minimization>
AUTHORS: Jianbo Chen, Mitchell Stern, Martin J. Wainwright, Michael I. Jordan
HIGHLIGHT: We propose a method for feature selection that employs kernel-based measures of independence to find a subset of covariates that is maximally predictive of the response.
- 665, TITLE: Convergence of Gradient EM on Multi-component Mixture of Gaussians
<https://papers.nips.cc/paper/7271-convergence-of-gradient-em-on-multi-component-mixture-of-gaussians>
AUTHORS: Bowei Yan, Mingzhang Yin, Purnamrita Sarkar
HIGHLIGHT: In this paper, we study convergence properties of the gradient variant of Expectation-Maximization algorithm~\cite{lange1995gradient} for Gaussian Mixture Models for arbitrary number of clusters and mixing coefficients.
- 666, TITLE: Real Time Image Saliency for Black Box Classifiers
<https://papers.nips.cc/paper/7272-real-time-image-saliency-for-black-box-classifiers>
AUTHORS: Piotr Dabkowski, Yarin Gal
HIGHLIGHT: In this work we develop a fast saliency detection method that can be applied to any differentiable image classifier.
- 667, TITLE: Houdini: Fooling Deep Structured Visual and Speech Recognition Models with Adversarial Examples
<https://papers.nips.cc/paper/7273-houdini-fooling-deep-structured-visual-and-speech-recognition-models-with-adversarial-examples>
AUTHORS: Moustapha M. Cisse, Yossi Adi, Natalia Neverova, Joseph Keshet
HIGHLIGHT: We introduce a novel flexible approach named Houdini for generating adversarial examples specifically tailored for the final performance measure of the task considered, be it combinatorial and non-decomposable.
- 668, TITLE: Efficient and Flexible Inference for Stochastic Systems
<https://papers.nips.cc/paper/7274-efficient-and-flexible-inference-for-stochastic-systems>
AUTHORS: Stefan Bauer, Nico S. Gorbach, Djordje Miladinovic, Joachim M. Buhmann
HIGHLIGHT: We provide a grid free and flexible algorithm offering parameter and state inference for stochastic systems and compare our approach based on variational approximations to state of the art methods showing significant advantages both in runtime and accuracy.
- 669, TITLE: When Cyclic Coordinate Descent Outperforms Randomized Coordinate Descent
<https://papers.nips.cc/paper/7275-when-cyclic-coordinate-descent-outperforms-randomized-coordinate-descent>
AUTHORS: Mert Gurbuzbalaban, Asuman Ozdaglar, Pablo A. Parrilo, Nuri Vanli
HIGHLIGHT: In this paper, we provide examples and more generally problem classes for which CCD (or CD with any deterministic order) is faster than RCD in terms of asymptotic worst-case convergence.
- 670, TITLE: Active Learning from Peers
<https://papers.nips.cc/paper/7276-active-learning-from-peers>
AUTHORS: Keerthiram Murugesan, Jaime Carbonell
HIGHLIGHT: The paper develops the new algorithm to exhibit this behavior and proves a theoretical mistake bound for the method compared to the best linear predictor in hindsight.
- 671, TITLE: Experimental Design for Learning Causal Graphs with Latent Variables
<https://papers.nips.cc/paper/7277-experimental-design-for-learning-causal-graphs-with-latent-variables>
AUTHORS: Murat Kocaoglu, Karthikeyan Shanmugam, Elias Bareinboim
HIGHLIGHT: We propose an efficient randomized algorithm that can learn the observable graph using $O(d \log^2 n)$ interventions where d is the degree of the graph.

- 672, TITLE: Learning to Model the Tail
<https://papers.nips.cc/paper/7278-learning-to-model-the-tail>
AUTHORS: Yu-Xiong Wang, Deva Ramanan, Martial Hebert
HIGHLIGHT: We describe an approach to learning from long-tailed, imbalanced datasets that are prevalent in real-world settings.
- 673, TITLE: Stochastic Mirror Descent in Variationally Coherent Optimization Problems
<https://papers.nips.cc/paper/7279-stochastic-mirror-descent-in-variationally-coherent-optimization-problems>
AUTHORS: Zhengyuan Zhou, Panayotis Mertikopoulos, Nicholas Bambos, Stephen Boyd, Peter W. Glynn
HIGHLIGHT: In this paper, we examine a class of non-convex stochastic optimization problems which we call variationally coherent, and which properly includes pseudo-/quasiconvex and star-convex optimization problems.
- 674, TITLE: On Separability of Loss Functions, and Revisiting Discriminative Vs Generative Models
<https://papers.nips.cc/paper/7280-on-separability-of-loss-functions-and-revisiting-discriminative-vs-generative-models>
AUTHORS: Adarsh Prasad, Alexandru Niculescu-Mizil, Pradeep K. Ravikumar
HIGHLIGHT: We revisit the classical analysis of generative vs discriminative models for general exponential families, and high-dimensional settings.
- 675, TITLE: Maxing and Ranking with Few Assumptions
<https://papers.nips.cc/paper/7281-maxing-and-ranking-with-few-assumptions>
AUTHORS: Moein Falahatgar, Yi Hao, Alon Orlitsky, Venkatadheeraj Pichapati, Vaishakh Ravindrakumar
HIGHLIGHT: With just one simple natural assumption: strong stochastic transitivity, we show that maxing can be performed with linearly many comparisons yet ranking requires quadratically many.
- 676, TITLE: On clustering network-valued data
<https://papers.nips.cc/paper/7282-on-clustering-network-valued-data>
AUTHORS: Soumendu Sundar Mukherjee, Purnamrita Sarkar, Lizhen Lin
HIGHLIGHT: When node correspondence is present, we cluster networks by summarizing a network by its graphon estimate, whereas when node correspondence is not present, we propose a novel solution for clustering such networks by associating a computationally feasible feature vector to each network based on trace of powers of the adjacency matrix.
- 677, TITLE: A General Framework for Robust Interactive Learning
<https://papers.nips.cc/paper/7283-a-general-framework-for-robust-interactive-learning>
AUTHORS: Ehsan Emamjomeh-Zadeh, David Kempe
HIGHLIGHT: We propose a general framework for interactively learning models, such as (binary or non-binary) classifiers, orderings/rankings of items, or clusterings of data points.
- 678, TITLE: Multi-view Matrix Factorization for Linear Dynamical System Estimation
<https://papers.nips.cc/paper/7284-multi-view-matrix-factorization-for-linear-dynamical-system-estimation>
AUTHORS: Mahdi Karami, Martha White, Dale Schuurmans, Csaba Szepesvari
HIGHLIGHT: In this paper, we instead reconsider likelihood maximization and develop an optimization based strategy for recovering the latent states and transition parameters.
- 679, TITLE: Mean Field Residual Networks: On the Edge of Chaos
<https://papers.nips.cc/paper/6879-mean-field-residual-networks-on-the-edge-of-chaos>
AUTHORS: Ge Yang, Samuel Schoenholz
HIGHLIGHT: We study randomly initialized residual networks using mean field theory and the theory of difference equations.