1, TITLE: Finding Biological Plausibility for Adversarially Robust Features via Metameric Tasks
https://openreview.net/forum?id=yeP_zx9vqNm
AUTHORS: Anne Harrington, Arturo Deza
HIGHLIGHT: We suggest that the representations learned by an Adversarially Trained Network are aligned with Human Peripheral Computation

2, TITLE: Task Relatedness-Based Generalization Bounds for Meta Learning
https://openreview.net/forum?id=A3HHaEdqAJL
AUTHORS: Jiechao Guan, Zhiwu Lu
HIGHLIGHT: In this paper, we propose to address this problem by defining a new notion of task relatedness according to the existence of the bijective transformation between two tasks.

3, TITLE: Compositional Attention: Disentangling Search and Retrieval
https://openreview.net/forum?id=IwJPj2MBcIa
AUTHORS: Sarthak Mittal, Sharath Chandra Raparthy, Irina Rish, Yoshua Bengio, Guillaume Lajoie
HIGHLIGHT: Recombining search and retrieval mechanisms of multi-head attention in a disentangled and flexible manner for better representational capacity and generalization.

4, TITLE: Constrained Policy Optimization via Bayesian World Models
https://openreview.net/forum?id=PRZoSmCinhf
AUTHORS: Yarden As, Ilnura Usmanova, Sebastian Curi, Andreas Krause
HIGHLIGHT: Solving constrained Markov decision processes with Bayesian model-based reinforcement learning.

5, TITLE: GNN-LM: Language Modeling based on Global Contexts via GNN
https://openreview.net/forum?id=BS49l-B5Bql
AUTHORS: Yuxian Meng, Shi Zong, Xiaoya Li, Xiaofei Sun, Tianwei Zhang, Fei Wu, Jiwei Li
HIGHLIGHT: Inspired by the notion that ``{\it to copy is easier than to memorize}``, in this work, we introduce GNN-LM, which extends vanilla neural language model (LM) by allowing to reference similar contexts in the entire training corpus.

6, TITLE: Unifying Likelihood-free Inference with Black-box Optimization and Beyond
https://openreview.net/forum?id=1HxTO6CTkz
AUTHORS: Dinghuai Zhang, Jie Fu, Yoshua Bengio, Aaron Courville
HIGHLIGHT: We propose a framework to unify likelihood-free inference and black-box sequence design and further propose novel sequence design algorithms based on the framework.

7, TITLE: PAC-Bayes Information Bottleneck
https://openreview.net/forum?id=iLHOIDsPv1P
AUTHORS: Zifeng Wang, Shao-Lun Huang, Ercan Engin Kuruoglu, Jimeng Sun, Xi Chen, Yefeng Zheng
HIGHLIGHT: We propose a novel PAC-Bayes bound guided information bottleneck for understanding and enhancing deep representation learning.

8, TITLE: Equivariant Transformers for Neural Network based Molecular Potentials
https://openreview.net/forum?id=zNHzqZ9wrRB
AUTHORS: Philipp Th?lke, Gianni De Fabritiis
HIGHLIGHT: We propose a novel equivariant Transformer architecture for the prediction of molecular potentials and provide insights into the molecular representation through extensive analysis of the model's attention weights.

9, TITLE: Scalable One-Pass Optimisation of High-Dimensional Weight-Update Hyperparameters by Implicit Differentiation
https://openreview.net/forum?id=hfU7Ka5cfrC
AUTHORS: Ross M Clarke, Elre Talea Oldewage, Jos? Miguel Hern?ndez-Lobato
HIGHLIGHT: We develop a gradient-based hyperparameter optimisation algorithm, applicable to a wide range of continuous hyperparameters, and scaling to large numbers of hyperparameters, without dramatically increasing training time from the non-HPO baseline.

10, TITLE: Distributional Decision Transformer for Hindsight Information Matching
https://openreview.net/forum?id=CAjxVodl_v
AUTHORS: Hiroki Furuta, Yutaka Matsuo, Shixiang Shane Gu
HIGHLIGHT: We generalize hindsight algorithms in RL, and propose Distributional Decision Transformer for information matching.

11, TITLE: Increasing the Cost of Model Extraction with Calibrated Proof of Work
https://openreview.net/forum?id=EAy7C1cgE1L
AUTHORS: Adam Dziedzic, Muhammad Ahmad Kaleem, Yu Shen Lu, Nicolas Papernot
HIGHLIGHT: We propose to make model extraction more difficult by requiring users to complete a callibrated proof-of-work before they can read predictions from a machine learning model exposed via a public API.

12, TITLE: Towards a Unified View of Parameter-Efficient Transfer Learning
https://openreview.net/forum?id=0RDcd5Axok
AUTHORS: Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, Graham Neubig
HIGHLIGHT: We propose a unified framework for several state-of-the-art parameter-efficient tuning methods,

13, TITLE: Towards Understanding the Data Dependency of Mixup-style Training
https://openreview.net/forum?id=ieNJYujcGDO
AUTHORS: Muthu Chidambaram, Xiang Wang, Yuzheng Hu, Chenwei Wu, Rong Ge
HIGHLIGHT: A theoretical analysis of data conditions under which mixup can perform worse, better, and identically when compared to empirical risk minimization.

14, TITLE: Escaping limit cycles: Global convergence for constrained nonconvex-nonconcave minimax problems
https://openreview.net/forum?id=2_vhkAMARk
AUTHORS: Thomas Pethick, Puya Latafat, Panos Patrinos, Olivier Fercoq, Volkan Cevher
HIGHLIGHT: Under weak MVI we introduce a new extragradient-type algorithm that avoids limit cycles

15, TITLE: Imbedding Deep Neural Networks
https://openreview.net/forum?id=yKIAXjkJc2F
AUTHORS: Andrew Corbett, Dmitry Kangin
HIGHLIGHT: Invariant imbedding solution for (Bolza) optimal control problem derived and proved to yield new architectures of imbedded deep neural networks.

16, TITLE: Progressive Distillation for Fast Sampling of Diffusion Models
https://openreview.net/forum?id=TIdIXIpzhoI
AUTHORS: Tim Salimans, Jonathan Ho
HIGHLIGHT: Diffusion models now need just 4 sampling steps to produce high quality samples.

17, TITLE: Spike-inspired rank coding for fast and accurate recurrent neural networks
https://openreview.net/forum?id=iMH1e5k7n3L
AUTHORS: Alan Jeffares, Qinghai Guo, Pontus Stenetorp, Timoleon Moraitis
HIGHLIGHT: Learning to infer fast in LSTMs inspired by SNNs, and applied in speech recognition

18, TITLE: Continual Learning with Filter Atom Swapping
https://openreview.net/forum?id=metRpM4Zrcb
AUTHORS: Zichen Miao, Ze Wang, Wei Chen, Qiang Qiu
HIGHLIGHT: In this paper, we first enforce a low-rank filter subspace by decomposing convolutional filters within each network layer over a small set of filter atoms.

19, TITLE: SHINE: SHaring the INverse Estimate from the forward pass for bi-level optimization and implicit models
https://openreview.net/forum?id=-ApAkox5mp
AUTHORS: Zaccharie Ramzi, Florian Mannel, Shaojie Bai, Jean-Luc Starck, Philippe Ciuciu, Thomas Moreau
HIGHLIGHT: Use the approximate Jacobian matrix computed in quasi-Newton methods to perform the inversion needed in the training of implicit models.

20, TITLE: Learning more skills through optimistic exploration
https://openreview.net/forum?id=cU8rknuhxc
AUTHORS: DJ Strouse, Kate Baumli, David Warde-Farley, Volodymyr Mnih, Steven Stenberg Hansen
HIGHLIGHT: Learn more skills by adding an information gain exploration bonus based on discriminator ensemble disagreement.

21, TITLE: Controlling Directions Orthogonal to a Classifier
https://openreview.net/forum?id=DIjCrlsu6Z
AUTHORS: Yilun Xu, Hao He, Tianxiao Shen, Tommi S. Jaakkola
HIGHLIGHT: We develop a notion of orthogonality in classifier, and the corresponding construction and utility.

22, TITLE: Revisiting Over-smoothing in BERT from the Perspective of Graph
https://openreview.net/forum?id=dUV91uaXm3
AUTHORS: Han Shi, JIAHUI GAO, Hang Xu, Xiaodan Liang, Zhenguo Li, Lingpeng Kong, Stephen M. S. Lee, James Kwok
HIGHLIGHT: We theoretically analyze the over-smoothing phenomenon of transformer-based models (e.g., BERT) and propose a novel hierarchical fusion strategy to alleviate it.

23, TITLE: Continual Learning with Recursive Gradient Optimization
https://openreview.net/forum?id=7YDLgf9_zgm
AUTHORS: Hao Liu, Huaping Liu
HIGHLIGHT: This paper proposes a novel method for continual learning in a fixed capacity network in the non-replay regime, which minimizes the loss on the current task while also minimizing an upper bound of loss increment on previous tasks.

24, TITLE: Improving Federated Learning Face Recognition via Privacy-Agnostic Clusters
https://openreview.net/forum?id=7l1IjZVddDW
AUTHORS: Qiang Meng, Feng Zhou, Hainan Ren, Tianshu Feng, Guochao Liu, Yuanqing Lin
HIGHLIGHT: To resolve the privacy-utility paradox, this work proposes PrivacyFace, a framework largely improves the federated learning face recognition via communicating auxiliary and privacy-agnostic information among clients.

25, TITLE: Pixelated Butterfly: Simple and Efficient Sparse training for Neural Network Models
https://openreview.net/forum?id=Nfl-iXa-y7R
AUTHORS: Beidi Chen, Tri Dao, Kaizhao Liang, Jiaming Yang, Zhao Song, Atri Rudra, Christopher Re
HIGHLIGHT: We propose a simple sparse training method, which can speed up model training in wall-clock time with no drop in accuracy.

26, TITLE: SGD Can Converge to Local Maxima
https://openreview.net/forum?id=9XhPLAjjRB
AUTHORS: Liu Ziyin, Botao Li, James B Simon, Masahito Ueda
HIGHLIGHT: We show that it can be common for SGD to converge to saddle points and maxima.

27, TITLE: Context-Aware Sparse Deep Coordination Graphs
https://openreview.net/forum?id=wQfgfb8VKTn
AUTHORS: Tonghan Wang, Liang Zeng, Weijun Dong, Qianlan Yang, Yang Yu, Chongjie Zhang
HIGHLIGHT: We propose a novel method for learning sparse coordination graphs that can be theoretically justified and can significantly reduce communication overhead and improve learning performance of deep coordination graphs.
To empirically evaluate our method, we present the Multi-Agent COordination (MACO) benchmark by collecting classic coordination problems in the literature, increasing their difficulty, and classifying them into different types.

28, TITLE: Learning Long-Term Reward Redistribution via Randomized Return Decomposition
https://openreview.net/forum?id=lpkGn3k2YdD
AUTHORS: Zhizhou Ren, Ruihan Guo, Yuan Zhou, Jian Peng
HIGHLIGHT: We propose randomized return decomposition, a novel reward redistribution algorithm, which establishes a surrogate optimization problem to scale up learning in long-horizon tasks.

29, TITLE: Fairness in Representation for Multilingual NLP: Insights from Controlled Experiments on Conditional Language Modeling
https://openreview.net/forum?id=-llS6TiOew
AUTHORS: Ada Wan
HIGHLIGHT: We investigate performance disparity in multilingual NLP with Transformer conditional LMs, and find, in the context of computing, morphological complexity to be a byproduct of word segmentation and disparity arising therefrom unwarranted.

30, TITLE: Understanding and Preventing Capacity Loss in Reinforcement Learning
https://openreview.net/forum?id=ZkC8wKoLbQ7
AUTHORS: Clare Lyle, Mark Rowland, Will Dabney
HIGHLIGHT: We show that RL agents experience representation collapse in sparse reward environments and propose an auxiliary task that prevents this from happening and outperforms the state of the art on the Atari benchmark.

31, TITLE: On the Importance of Firth Bias Reduction in Few-Shot Classification
https://openreview.net/forum?id=DNRADop4ksB

AUTHORS: Saba Ghaffari, Ehsan Saleh, David Forsyth, Yu-Xiong Wang
HIGHLIGHT: In this work, we demonstrate the effectiveness of Firth bias reduction in few-shot classification.

32, TITLE: RotoGrad: Gradient Homogenization in Multitask Learning
https://openreview.net/forum?id=T8wHz4rnuGL
AUTHORS: Adri?n Javaloy, Isabel Valera
HIGHLIGHT: We propose an algorithm to simultaneously homogenize gradient magnitudes and directions across tasks in MTL.

33, TITLE: Learning the Dynamics of Physical Systems from Sparse Observations with Finite Element Networks
https://openreview.net/forum?id=HFmAukZ-k-2
AUTHORS: Marten Lienen, Stephan G?nnemann
HIGHLIGHT: A continuous-time graph neural network model for spatio-temporal forecasting that can structurally incorporate prior knowledge

34, TITLE: Responsible Disclosure of Generative Models Using Scalable Fingerprinting
https://openreview.net/forum?id=sOK-zS6WHB
AUTHORS: Ning Yu, Vladislav Skripniuk, Dingfan Chen, Larry S. Davis, Mario Fritz
HIGHLIGHT: Our work enables a responsible disclosure of generative models, that allows model inventors to fingerprint their models, so that the generated samples containing a fingerprint can be accurately detected and attributed to a source.

35, TITLE: How to Robustify Black-Box ML Models? A Zeroth-Order Optimization Perspective
https://openreview.net/forum?id=W9G_ImpHlQd
AUTHORS: Yimeng Zhang, Yuguang Yao, Jinghan Jia, Jinfeng Yi, Mingyi Hong, Shiyu Chang, Sijia Liu
HIGHLIGHT: We propose a general notion of defensive operation that can be applied to black-box models, and design it through the lens of denoised smoothing (DS), a first-order (FO) certified defense technique.

36, TITLE: Learning Multimodal VAEs through Mutual Supervision
https://openreview.net/forum?id=1xXvPrAshao
AUTHORS: Tom Joy, Yuge Shi, Philip Torr, Tom Rainforth, Sebastian M Schmon, Siddharth N
HIGHLIGHT: Here we re-purpose semi-supervised VAEs to leverage mutual supervision between encoding distributions, allowing us to learn multi-modal VAEs with partially obsereved data.

37, TITLE: Reinforcement Learning under a Multi-agent Predictive State Representation Model: Method and Theory
https://openreview.net/forum?id=PLDOnFoVm4
AUTHORS: Zhi Zhang, Zhuoran Yang, Han Liu, Pratap Tokekar, Furong Huang
HIGHLIGHT: We propose a new algorithm for MARL under a multi-agent predictive state representation model, where we incorporate a dynamic interaction graph; we provide the theoretical guarantees of our model and run various experiments to support our algorithm.

38, TITLE: Perceiver IO: A General Architecture for Structured Inputs & Outputs
https://openreview.net/forum?id=fILj7WpI-g
AUTHORS: Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, Olivier J Henaff, Matthew Botvinick, Andrew Zisserman, Oriol Vinyals, Joao Carreira
HIGHLIGHT: We propose Perceiver IO, a general-purpose architecture that handles data from arbitrary settings while scaling linearly with the size of inputs and outputs.

39, TITLE: Natural Posterior Network: Deep Bayesian Predictive Uncertainty for Exponential Family Distributions
https://openreview.net/forum?id=tV3N0DWMxCg
AUTHORS: Bertrand Charpentier, Oliver Borchert, Daniel Z?gner, Simon Geisler, Stephan G?nnemann
HIGHLIGHT: In this work, we propose the Natural Posterior Network (NatPN) for fast and high-quality uncertainty estimation for any task where the target distribution belongs to the exponential family.

40, TITLE: Spanning Tree-based Graph Generation for Molecules
https://openreview.net/forum?id=w60btE_8T2m
AUTHORS: Sungsoo Ahn, Binghong Chen, Tianzhe Wang, Le Song
HIGHLIGHT: We propose a new molecular graph generative model based on compact tree constructive operators.

41, TITLE: Meta Discovery: Learning to Discover Novel Classes given Very Limited Data

https://openreview.net/forum?id=MEpKGLsY8f
AUTHORS: Haoang Chi, Feng Liu, Wenjing Yang, Long Lan, Tongliang Liu, Bo Han, Gang Niu, Mingyuan Zhou, Masashi Sugiyama
HIGHLIGHT: In this paper, we demystify assumptions behind L2DNC and find that high-level semantic features should be shared among the seen and unseen classes.

42, TITLE: Training invariances and the low-rank phenomenon: beyond linear networks
https://openreview.net/forum?id=XEW8CQgArno
AUTHORS: Thien Le, Stefanie Jegelka
HIGHLIGHT: We extend theoretical results regarding the low-rank bias of deep linear neural networks trained with gradient-based algorithm to non-linear architectures, reflecting empirical results in the literature.

43, TITLE: On Improving Adversarial Transferability of Vision Transformers
https://openreview.net/forum?id=D6nH3719vZy
AUTHORS: Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Khan, Fatih Porikli
HIGHLIGHT: Novel approach to improve transferability of adversarial perturbations found in vision transformers via self-ensemble and token refinement.

44, TITLE: Scarf: Self-Supervised Contrastive Learning using Random Feature Corruption
https://openreview.net/forum?id=CuV_qYkmKb3
AUTHORS: Dara Bahri, Heinrich Jiang, Yi Tay, Donald Metzler
HIGHLIGHT: Scarf is a self-supervised, contrastive pre-training method for neural networks applied to tabular classification tasks that boosts performance, even when labeled data is limited or noisy.

45, TITLE: 8-bit Optimizers via Block-wise Quantization
https://openreview.net/forum?id=shpkpVXzo3h
AUTHORS: Tim Dettmers, Mike Lewis, Sam Shleifer, Luke Zettlemoyer
HIGHLIGHT: We develop 8-bit optimizers reduce the memory footprint of training and maintain 32-bit optimizer performance across NLP/CV benchmarks.

46, TITLE: Superclass-Conditional Gaussian Mixture Model For Learning Fine-Grained Embeddings
https://openreview.net/forum?id=vds4SNooOe
AUTHORS: Jingchao Ni, Wei Cheng, Zhengzhang Chen, Takayoshi Asakura, Tomoya Soma, Sho Kato, Haifeng Chen
HIGHLIGHT: We propose a training framework characterized by a novel superclass conditional Gaussian mixture (SCGM) based generative model for learning fine-grained representations for cross-granularity adaptation.

47, TITLE: Pessimistic Bootstrapping for Uncertainty-Driven Offline Reinforcement Learning
https://openreview.net/forum?id=Y4cs1Z3HnqL
AUTHORS: Chenjia Bai, Lingxiao Wang, Zhuoran Yang, Zhi-Hong Deng, Animesh Garg, Peng Liu, Zhaoran Wang
HIGHLIGHT: We propose pessimistic bootstrapping as a purely uncertainty-driven algorithm for offline Reinforcement Learning.

48, TITLE: Iterative Refinement Graph Neural Network for Antibody Sequence-Structure Co-design
https://openreview.net/forum?id=LI2bhrE_2A
AUTHORS: Wengong Jin, Jeremy Wohlwend, Regina Barzilay, Tommi S. Jaakkola
HIGHLIGHT: We propose a new graph-based generative model for antibody design

49, TITLE: Latent Variable Sequential Set Transformers for Joint Multi-Agent Motion Prediction
https://openreview.net/forum?id=Dup_dDqkZC5
AUTHORS: Roger Girgis, Florian Golemo, Felipe Codevilla, Martin Weiss, Jim Aldon D'Souza, Samira Ebrahimi Kahou, Felix Heide, Christopher Pal
HIGHLIGHT: New Transformer-based architecture for socially consistent motion forecasting. Achieves SotA performance on NuScenes at a fraction of the compute of competing methods.

50, TITLE: Autoregressive Quantile Flows for Predictive Uncertainty Estimation
https://openreview.net/forum?id=z1-I6rOKv1S
AUTHORS: Phillip Si, Volodymyr Kuleshov, Allan Bishop
HIGHLIGHT: Using Quantile Flows for Predictive and Generative Data Modeling and Generation

51, TITLE: Analytic-DPM: an Analytic Estimate of the Optimal Reverse Variance in Diffusion Probabilistic Models

https://openreview.net/forum?id=0xiJLKH-ufZ
AUTHORS: Fan Bao, Chongxuan Li, Jun Zhu, Bo Zhang
HIGHLIGHT: We propose an analytic framework of estimating the optimal reverse variance in DPMs.

52, TITLE: Coordination Among Neural Modules Through a Shared Global Workspace
https://openreview.net/forum?id=XzTtHjgPDsT
AUTHORS: Anirudh Goyal, Aniket Rajiv Didolkar, Alex Lamb, Kartikeya Badola, Nan Rosemary Ke, Nasim Rahaman, Jonathan Binas, Charles Blundell, Michael Curtis Mozer, Yoshua Bengio
HIGHLIGHT: communication among different specialist using a shared workspace allowing higher order interactions

53, TITLE: Weighted Training for Cross-Task Learning
https://openreview.net/forum?id=ltM1RMZntpu
AUTHORS: Shuxiao Chen, Koby Crammer, Hangfeng He, Dan Roth, Weijie J Su
HIGHLIGHT: We introduce a weighted training algorithm for cross-task learning based on minimizing a representation-based task distance between the source and target tasks.

54, TITLE: Poisoning and Backdooring Contrastive Learning
https://openreview.net/forum?id=iC4UHbQ01Mp
AUTHORS: Nicholas Carlini, Andreas Terzis
HIGHLIGHT: We argue poisoning and backdooring attacks are a serious threat to multimodal contrastive classifiers, because they are explicitly designed to be trained on uncurated datasets from the Internet.

55, TITLE: Generative Planning for Temporally Coordinated Exploration in Reinforcement Learning
https://openreview.net/forum?id=YZHES8wIdE
AUTHORS: Haichao Zhang, Wei Xu, Haonan Yu
HIGHLIGHT: Temporally coordinated exploration in reinforcement learning using Generative Planning Method.

56, TITLE: SphereFace2: Binary Classification is All You Need for Deep Face Recognition
https://openreview.net/forum?id=l3SDgUh7qZO
AUTHORS: Yandong Wen, Weiyang Liu, Adrian Weller, Bhiksha Raj, Rita Singh
HIGHLIGHT: A novel deep face recognition framework

57, TITLE: Possibility Before Utility: Learning And Using Hierarchical Affordances
https://openreview.net/forum?id=7b4zxUnrO2N
AUTHORS: Robby Costales, Shariq Iqbal, Fei Sha
HIGHLIGHT: We introduce a method that achieves superior performance in complex hierarchical tasks by utilizing a notion of subtask dependency grounded in the present state.

58, TITLE: IntSGD: Adaptive Floatless Compression of Stochastic Gradients
https://openreview.net/forum?id=pFyXqxChZc
AUTHORS: Konstantin Mishchenko, Bokun Wang, Dmitry Kovalev, Peter Richt?rik
HIGHLIGHT: We propose the provably convergent and computationally cheap IntSGD algorithm for efficient distributed machine learning.

59, TITLE: Score-Based Generative Modeling with Critically-Damped Langevin Diffusion
https://openreview.net/forum?id=CzceR82CYc
AUTHORS: Tim Dockhorn, Arash Vahdat, Karsten Kreis
HIGHLIGHT: In this work, we propose a novel diffusion process ideally suited for score-based generative models and provide new insights into score-based denoising diffusion models.

60, TITLE: Revisiting Design Choices in Offline Model Based Reinforcement Learning
https://openreview.net/forum?id=zz9hXVhf40
AUTHORS: Cong Lu, Philip Ball, Jack Parker-Holder, Michael Osborne, Stephen J. Roberts
HIGHLIGHT: In this paper, we compare these heuristics, and design novel protocols to investigate their interaction with other hyperparameters, such as the number of models, or imaginary rollout horizon.

61, TITLE: Value Gradient weighted Model-Based Reinforcement Learning
https://openreview.net/forum?id=4-D6CZkRXxI
AUTHORS: Claas A Voelcker, Victor Liao, Animesh Garg, Amir-massoud Farahmand

HIGHLIGHT: We propose the Value-gradient weighted Model loss, a method for value-aware model learning in challenging settings, such as small model capacity and the presence of distracting state dimensions.

62, TITLE: Memorizing Transformers
https://openreview.net/forum?id=TrjbxzRcnf-
AUTHORS: Yuhuai Wu, Markus Norman Rabe, DeLesley Hutchins, Christian Szegedy
HIGHLIGHT: We propose to use an external memory module to allow instant utilization of newly acquired knowledge.

63, TITLE: Scaling Laws for Neural Machine Translation
https://openreview.net/forum?id=hR_SMu8cxCV
AUTHORS: Behrooz Ghorbani, Orhan Firat, Markus Freitag, Ankur Bapna, Maxim Krikun, Xavier Garcia, Ciprian Chelba, Colin Cherry
HIGHLIGHT: We provide (model) scaling laws for neural machine translation.

64, TITLE: Multi-Stage Episodic Control for Strategic Exploration in Text Games
https://openreview.net/forum?id=Ek7PSN7Y77z
AUTHORS: Jens Tuyls, Shunyu Yao, Sham M. Kakade, Karthik R Narasimhan
HIGHLIGHT: We propose a multi-stage approach to playing text games that improves the score on Zork1 from around 40 to 103.

65, TITLE: Explanations of Black-Box Models based on Directional Feature Interactions
https://openreview.net/forum?id=45Mr7LeKR9
AUTHORS: Aria Masoomi, Davin Hill, Zhonghui Xu, Craig P Hersh, Edwin K. Silverman, Peter J. Castaldi, Stratis Ioannidis, Jennifer Dy
HIGHLIGHT: We introduce a bivariate explainer to explain directional feature interactions in black box models.

66, TITLE: Programmatic Reinforcement Learning without Oracles
https://openreview.net/forum?id=6Tk2noBdvxt
AUTHORS: Wenjie Qiu, He Zhu
HIGHLIGHT: We present a differentiable program architecture search framework to synthesize interpretable, generalizable, and compositional programs for controlling reinforcement learning applications.

67, TITLE: Variational methods for simulation-based inference
https://openreview.net/forum?id=kZ0UYdhqkNY
AUTHORS: Manuel Gl?ckler, Michael Deistler, Jakob H. Macke
HIGHLIGHT: We combine likelihood-estimation with variational inference to achieve a scalable approach for simulation-based inference.

68, TITLE: Learning transferable motor skills with hierarchical latent mixture policies
https://openreview.net/forum?id=qTHBE7E9iej
AUTHORS: Dushyant Rao, Fereshteh Sadeghi, Leonard Hasenclever, Markus Wulfmeier, Martina Zambelli, Giulia Vezzani, Dhruva Tirumala, Yusuf Aytar, Josh Merel, Nicolas Heess, raia hadsell
HIGHLIGHT: An approach to learn reusable and transferable skills from data via a hierarchical latent mixture policy, which can significantly improve sample efficiency and asymptotic performance on downstream RL tasks

69, TITLE: On the relation between statistical learning and perceptual distances
https://openreview.net/forum?id=zXM0b4hi5_B
AUTHORS: Alexander Hepburn, Valero Laparra, Raul Santos-Rodriguez, Johannes Ball?, Jesus Malo
HIGHLIGHT: In this paper, we aim to unravel the non-trivial relationships between the probability distribution of the data, perceptual distances, and unsupervised machine learning.

70, TITLE: Half-Inverse Gradients for Physical Deep Learning
https://openreview.net/forum?id=HTx7vrlLBEj
AUTHORS: Patrick Schnell, Philipp Holl, Nils Thuerey
HIGHLIGHT: By proposing a novel Jacobian-based optimizer, we question the current practice of using the state-of-the-art gradient-based methods for the optimization of neural networks with physics objectives.

71, TITLE: NODE-GAM: Neural Generalized Additive Model for Interpretable Deep Learning
https://openreview.net/forum?id=g8NJR6fCCl8
AUTHORS: Chun-Hao Chang, Rich Caruana, Anna Goldenberg

HIGHLIGHT: We develop a deep-learning version of Generalized Additive Model (GAM) and GA2M that is accurate, scalable and interpretable.

72, TITLE: Graph-Augmented Normalizing Flows for Anomaly Detection of Multiple Time Series
https://openreview.net/forum?id=45L_dgP48Vd
AUTHORS: Enyan Dai, Jie Chen
HIGHLIGHT: We hypothesize that anomalies occur in low density regions of a distribution and explore the use of normalizing flows for unsupervised anomaly detection, because of their superior quality in density estimation.

73, TITLE: On Predicting Generalization using GANs
https://openreview.net/forum?id=eW5R4Cek6y6
AUTHORS: Yi Zhang, Arushi Gupta, Nikunj Saunshi, Sanjeev Arora
HIGHLIGHT: The current paper investigates a simple idea: can test error be predicted using {\em synthetic data,} produced using a Generative Adversarial Network (GAN) that was trained on the same training dataset?

74, TITLE: Optimal Transport for Causal Discovery
https://openreview.net/forum?id=qwBK94cP1y
AUTHORS: Ruibo Tu, Kun Zhang, Hedvig Kjellstrom, Cheng Zhang
HIGHLIGHT: In this paper, we provide a novel dynamical-system view of FCMs and propose a new framework for identifying causal direction in the bivariate case.

75, TITLE: On Bridging Generic and Personalized Federated Learning for Image Classification
https://openreview.net/forum?id=I1hQbx10Kxn
AUTHORS: Hong-You Chen, Wei-Lun Chao
HIGHLIGHT: Concretely, we propose a novel federated learning framework that explicitly decouples a model's dual duties with two prediction tasks.

76, TITLE: Compositional Training for End-to-End Deep AUC Maximization
https://openreview.net/forum?id=gPvB4pdu_Z
AUTHORS: Zhuoning Yuan, Zhishuai Guo, Nitesh Chawla, Tianbao Yang
HIGHLIGHT: We propose a novel end-to-end training framework with a provable stochastic algorithm for deep AUC maximization.

77, TITLE: Learning meta-features for AutoML
https://openreview.net/forum?id=DTkEfj0Ygb8
AUTHORS: Herilalaina Rakotoarison, Louisot Milijaona, Andry RASOANAIVO, Michele Sebag, Marc Schoenauer
HIGHLIGHT: We propose a novel approach to learn dataset meta-features for AutoML.

78, TITLE: $\mathrm{SO}(2)$-Equivariant Reinforcement Learning
https://openreview.net/forum?id=7F9cOhdvfk_
AUTHORS: Dian Wang, Robin Walters, Robert Platt
HIGHLIGHT: This paper proposes equivariant DQN and equivariant SAC that significantly improve the sample efficiency of RL in robotic manipulation.

79, TITLE: Byzantine-Robust Learning on Heterogeneous Datasets via Bucketing
https://openreview.net/forum?id=jXKKDEi5vJt
AUTHORS: Sai Praneeth Karimireddy, Lie He, Martin Jaggi
HIGHLIGHT: Byzantine-robust distributed learning with heterogeneous data distribution

80, TITLE: Learnability of convolutional neural networks for infinite dimensional input via mixed and anisotropic smoothness
https://openreview.net/forum?id=dgxFTxuJ50e
AUTHORS: Sho Okumoto, Taiji Suzuki
HIGHLIGHT: In this paper, we investigate the approximation and estimation errors of the (dilated) convolutional neural networks when the input is infinite dimensional.

81, TITLE: Omni-Dimensional Dynamic Convolution
https://openreview.net/forum?id=DmpCfq6Mg39
AUTHORS: Chao Li, Aojun Zhou, Anbang Yao

HIGHLIGHT: This paper presents Omni-dimensional Dynamic Convolution (ODConv) to advance the research in dynamic convolution.

82, TITLE: COptiDICE: Offline Constrained Reinforcement Learning via Stationary Distribution Correction Estimation
https://openreview.net/forum?id=FLA55mBee6Q
AUTHORS: Jongmin Lee, Cosmin Paduraru, Daniel J Mankowitz, Nicolas Heess, Doina Precup, Kee-Eung Kim, Arthur Guez
HIGHLIGHT: We present an offline constrained RL algorithm, which estimates the stationary distribution corrections of the optimal policy with respect to returns, while constraining the cost upper bound.

83, TITLE: Planning in Stochastic Environments with a Learned Model
https://openreview.net/forum?id=X6D9bAHhBQ1
AUTHORS: Ioannis Antonoglou, Julian Schrittwieser, Sherjil Ozair, Thomas K Hubert, David Silver
HIGHLIGHT: Specifically, we introduce a new algorithm, Stochastic MuZero, that learns a stochastic model incorporating afterstates, and uses this model to perform a stochastic tree search.

84, TITLE: Learning Optimal Conformal Classifiers
https://openreview.net/forum?id=t8O-4LKFVx
AUTHORS: David Stutz, Krishnamurthy Dj Dvijotham, Ali Taylan Cemgil, Arnaud Doucet
HIGHLIGHT: Conformal training allows to train classifier and conformal predictor end-to-end, optimizing average confidence set size (inefficiency) or other application-specific losses defined on confidence sets.

85, TITLE: Learning Causal Relationships from Conditional Moment Restrictions by Importance Weighting
https://openreview.net/forum?id=7twQI5VnC8
AUTHORS: Masahiro Kato, Masaaki Imaizumi, Kenichiro McAlinn, Shota Yasui, Haruo Kakehi
HIGHLIGHT: Learning causal relationships under conditional moment restrictions by importance weighting using the conditional density ratio function.

86, TITLE: Finite-Time Convergence and Sample Complexity of Multi-Agent Actor-Critic Reinforcement Learning with Average Reward
https://openreview.net/forum?id=04pGUg0-pdZ
AUTHORS: FNU Hairi, Jia Liu, Songtao Lu
HIGHLIGHT: In this paper, we establish the first finite-time convergence result of the actor-critic algorithm for fully decentralized multi-agent reinforcement learning (MARL) problems with average reward.

87, TITLE: Looking Back on Learned Experiences For Class/task Incremental Learning
https://openreview.net/forum?id=RxplU3vmBx
AUTHORS: Mozhgan PourKeshavarzi, Guoying Zhao, Mohammad Sabokrou
HIGHLIGHT: In this paper, we shed light on an on-call transfer set to provide past experiences whenever a new task arises in the data stream.

88, TITLE: TRGP: Trust Region Gradient Projection for Continual Learning
https://openreview.net/forum?id=iEvAf8i6JjO
AUTHORS: Sen Lin, Li Yang, Deliang Fan, Junshan Zhang
HIGHLIGHT: We propose a novel continual learning approach to facilitate the forward knowledge transfer, based on an efficient characterization of task correlation using a novel notion of 'trust region'.

89, TITLE: Analyzing and Improving the Optimization Landscape of Noise-Contrastive Estimation
https://openreview.net/forum?id=eBS-3YiaIL-
AUTHORS: Bingbin Liu, Elan Rosenfeld, Pradeep Kumar Ravikumar, Andrej Risteski
HIGHLIGHT: This work theoretically explains the difficulty of optimizing the NCE loss when the noise distribution is poor, and provides a provably efficient solution consisting of normalized gradient descent (NGD) combined with the proposed \emph{eNCE} loss.

90, TITLE: Understanding Domain Randomization for Sim-to-real Transfer
https://openreview.net/forum?id=T8vZHIRTrY
AUTHORS: Xiaoyu Chen, Jiachen Hu, Chi Jin, Lihong Li, Liwei Wang
HIGHLIGHT: We propose theoretical frameworks for sim-to-real transfer and domain randomization, and provide bounds on the sub-optimality gap of the policy returned by domain randomization.

91, TITLE: Lossless Compression with Probabilistic Circuits
https://openreview.net/forum?id=X_hByk2-5je
AUTHORS: Anji Liu, Stephan Mandt, Guy Van den Broeck
HIGHLIGHT: To overcome such problems, we establish a new class of tractable lossless compression models that permit efficient encoding and decoding: Probabilistic Circuits (PCs).

92, TITLE: Amortized Tree Generation for Bottom-up Synthesis Planning and Synthesizable Molecular Design
https://openreview.net/forum?id=FRxhHdnxt1
AUTHORS: Wenhao Gao, Roc?o Mercado, Connor W. Coley
HIGHLIGHT: We propose a model that address synthesis planning and synthesizable molecular design simultaneously.

93, TITLE: Wiring Up Vision: Minimizing Supervised Synaptic Updates Needed to Produce a Primate Ventral Stream
https://openreview.net/forum?id=g1SzIRLQXMM
AUTHORS: Franziska Geiger, Martin Schrimpf, Tiago Marques, James J. DiCarlo
HIGHLIGHT: We develop biologically-motivated initialization and training procedures to train models with 200x fewer synaptic updates (epochs x labeled images x weights) while maintaining 80% of brain predictivity on a set of neural and behavioral benchmarks.

94, TITLE: Hybrid Local SGD for Federated Learning with Heterogeneous Communications
https://openreview.net/forum?id=H0oaWl6THa
AUTHORS: Yuanxiong Guo, Ying Sun, Rui Hu, Yanmin Gong
HIGHLIGHT: In light of both device-to-device (D2D) and device-to-server (D2E) cooperation opportunities in modern communication networks, this paper proposes a new federated optimization algorithm dubbed hybrid local SGD (HL-SGD) in FL settings where devices are grouped into a set of disjoint clusters with high D2D communication bandwidth.

95, TITLE: Sample Efficient Deep Reinforcement Learning via Uncertainty Estimation
https://openreview.net/forum?id=vrW3tvDfOJQ
AUTHORS: Vincent Mai, Kaustubh Mani, Liam Paull
HIGHLIGHT: The sample efficiency and performance of model-free DRL is improved by estimating the predictive uncertainty of the targets using probabilistic ensembles and down-weighting the uncertain samples using batch inverse-variance weighting.

96, TITLE: Understanding the Role of Self Attention for Efficient Speech Recognition
https://openreview.net/forum?id=AvcfxqRy4Y
AUTHORS: Kyuhong Shim, Jungwook Choi, Wonyong Sung
HIGHLIGHT: We analyze the role of self attention in Transformer-based speech recognition and present a practical technique to design a model that accelerates the inference and improve the performance.

97, TITLE: Boosting Randomized Smoothing with Variance Reduced Classifiers
https://openreview.net/forum?id=mHu2vIds_-b
AUTHORS: Mikl?s Z. Horv?th, Mark Niklas Mueller, Marc Fischer, Martin Vechev
HIGHLIGHT: We show -- theoretically and empirically -- that ensembles reduce variance under randomized smoothing, yielding higher certified accuracy, leading to a new state-of-the-art on CIFAR-10 and ImageNet.

98, TITLE: Learning Altruistic Behaviours in Reinforcement Learning without External Rewards
https://openreview.net/forum?id=KxbhdyiPHE
AUTHORS: Tim Franzmeyer, Mateusz Malinowski, Joao F. Henriques
HIGHLIGHT: We propose and investigate unsupervised training of agents to behave altruistically towards others by actively maximizing others' choice.

99, TITLE: EViT: Expediting Vision Transformers via Token Reorganizations
https://openreview.net/forum?id=BjyvwnXXVn_
AUTHORS: Youwei Liang, Chongjian GE, Zhan Tong, Yibing Song, Jue Wang, Pengtao Xie
HIGHLIGHT: We propose to reorganize attentive tokens in Vision Transformers to expedite inference speed.

100, TITLE: Learning Vision-Guided Quadrupedal Locomotion End-to-End with Cross-Modal Transformers
https://openreview.net/forum?id=nhnJ3oo6AB
AUTHORS: Ruihan Yang, Minghao Zhang, Nicklas Hansen, Huazhe Xu, Xiaolong Wang
HIGHLIGHT: We introduce a novel end-to-end Reinforcement Learning approach called LocoTransformer, leveraging both visual inputs and proprioceptive states, for locomotion control in both simulation and with real robots.

101, TITLE: Adversarial Support Alignment
https://openreview.net/forum?id=26gKg6x-ie
AUTHORS: Shangyuan Tong, Timur Garipov, Yang Zhang, Shiyu Chang, Tommi S. Jaakkola
HIGHLIGHT: We study the problem of aligning the supports of distributions.

102, TITLE: Sampling with Mirrored Stein Operators
https://openreview.net/forum?id=eMudnJsb1T5
AUTHORS: Jiaxin Shi, Chang Liu, Lester Mackey
HIGHLIGHT: We introduce multi-particle generalization of mirror descent for sampling in constrained domains and non-Euclidean geometries.

103, TITLE: DR3: Value-Based Deep Reinforcement Learning Requires Explicit Regularization
https://openreview.net/forum?id=POvMvLi91f
AUTHORS: Aviral Kumar, Rishabh Agarwal, Tengyu Ma, Aaron Courville, George Tucker, Sergey Levine
HIGHLIGHT: We show that implicit regularization effects can lead to poor performance in value-based offline RL and propose an explicit regularizer to mitigate these effects.

104, TITLE: Learning Hierarchical Structures with Differentiable Nondeterministic Stacks
https://openreview.net/forum?id=5LXw_QplBiF
AUTHORS: Brian DuSell, David Chiang
HIGHLIGHT: We present a new stack-augmented RNN with strong results on CFL language modeling tasks.

105, TITLE: Exploring the Limits of Large Scale Pre-training
https://openreview.net/forum?id=V3C8p78sDa
AUTHORS: Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, Hanie Sedghi
HIGHLIGHT: We perform a systematic investigation of limits of large scale pre-training for few-shot and transfer learning in image recognition with a wide range of downstream tasks.

106, TITLE: Self-supervised Learning is More Robust to Dataset Imbalance
https://openreview.net/forum?id=4AZz9osqrar
AUTHORS: Hong Liu, Jeff Z. HaoChen, Adrien Gaidon, Tengyu Ma
HIGHLIGHT: We show that self-supervised pre-training yields representations more robust to dataset imbalance, because it captures more diverse features from the frequent classes, and can be improved further by re-weighting regularization.

107, TITLE: Probabilistic Implicit Scene Completion
https://openreview.net/forum?id=BnQhMqDfcKG
AUTHORS: Dongsu Zhang, Changwoon Choi, Inbum Park, Young Min Kim
HIGHLIGHT: We propose a scalable generative model for multi-modal completion of 3D scenes in implicit representation.

108, TITLE: Near-Optimal Reward-Free Exploration for Linear Mixture MDPs with Plug-in Solver
https://openreview.net/forum?id=SidzxAb9k30
AUTHORS: Xiaoyu Chen, Jiachen Hu, Lin Yang, Liwei Wang
HIGHLIGHT: We propose near-optimal exploration algorithms for reward-free exploration with plug-in solver.

109, TITLE: MT3: Multi-Task Multitrack Music Transcription
https://openreview.net/forum?id=iMSjopcOn0p
AUTHORS: Joshua P Gardner, Ian Simon, Ethan Manilow, Curtis Hawthorne, Jesse Engel
HIGHLIGHT: Unified framework for music transcription, jointly training a single model on six multi-instrument datasets and establishing a new SOTA for low-resource music transcription.

110, TITLE: Improved deterministic l2 robustness on CIFAR-10 and CIFAR-100
https://openreview.net/forum?id=tD7eCtaSkR
AUTHORS: Sahil Singla, Surbhi Singla, Soheil Feizi
HIGHLIGHT: Improving provable robustness of 1 Lipschitz CNNs by relaxing orthogonalization of last layer, certificate regularization and a novel activation function.

111, TITLE: Contact Points Discovery for Soft-Body Manipulations with Differentiable Physics
https://openreview.net/forum?id=mmUA7_O9mjY
AUTHORS: Sizhe Li, Zhiao Huang, Tao Du, Hao Su, Joshua B. Tenenbaum, Chuang Gan

HIGHLIGHT: We propose a contact pose discovery method that guides the stand-alone differentiable physics solver to complete various soft-body manipulation tasks.

112, TITLE: AdaRL: What, Where, and How to Adapt in Transfer Reinforcement Learning
https://openreview.net/forum?id=8H5bpVwvt5
AUTHORS: Biwei Huang, Fan Feng, Chaochao Lu, Sara Magliacane, Kun Zhang
HIGHLIGHT: Efficient policy adaptation across domains by learning a parsimonious graphical representation that encodes changes in a compact way.

113, TITLE: Learning Pruning-Friendly Networks via Frank-Wolfe: One-Shot, Any-Sparsity, And No Retraining
https://openreview.net/forum?id=O1DEtITim__
AUTHORS: Lu Miao, Xiaolong Luo, Tianlong Chen, Wuyang Chen, Dong Liu, Zhangyang Wang
HIGHLIGHT: We propose a novel and state-of-the-art one-shot pruning method, which can generate sparse networks at any pruning ratio in one pruning and without any retraining.

114, TITLE: Reinforcement Learning with Sparse Rewards using Guidance from Offline Demonstration
https://openreview.net/forum?id=YJ1WzgMVsMt
AUTHORS: Desik Rengarajan, Gargi Vaidya, Akshay Sarvesh, Dileep Kalathil, Srinivas Shakkottai
HIGHLIGHT: Reinforcement learning in sparse reward environments using offline guidance.

115, TITLE: The Inductive Bias of In-Context Learning: Rethinking Pretraining Example Design
https://openreview.net/forum?id=lnEaqbTJIRz
AUTHORS: Yoav Levine, Noam Wies, Daniel Jannai, Dan Navon, Yedid Hoshen, Amnon Shashua
HIGHLIGHT: We prove that pertained LMs model stronger dependencies between sentences that were shown in same training example, thus indicating benefits of better informed "pretraining example design"

116, TITLE: Emergent Communication at Scale
https://openreview.net/forum?id=AUGBfDIV9rL
AUTHORS: Rahma Chaabouni, Florian Strub, Florent Altch?, Eugene Tarassov, Corentin Tallec, Elnaz Davoodi, Kory Wallace Mathewson, Olivier Tieleman, Angeliki Lazaridou, Bilal Piot
HIGHLIGHT: This work argues the importance of scaling up the emergent communication framework and investigates the impact of three scaling up aspects, namely the dataset, task complexity, and population size.
We provide a first set of results for large populations solving complex tasks on realistic large-scale datasets, as well as an easy-to-use codebase to enable further experimentation.

117, TITLE: Long Expressive Memory for Sequence Modeling
https://openreview.net/forum?id=vwj6aUeocyf
AUTHORS: T. Konstantin Rusch, Siddhartha Mishra, N. Benjamin Erichson, Michael W. Mahoney
HIGHLIGHT: A novel method for sequence modeling based on multiscale ODEs that is provably able to learn very long-term dependencies while being sufficiently expressive to outperform state-of-the-art recurrent sequence models.

118, TITLE: When should agents explore?
https://openreview.net/forum?id=dEwfxt14bca
AUTHORS: Miruna Pislar, David Szepesvari, Georg Ostrovski, Diana L Borsa, Tom Schaul
HIGHLIGHT: A fresh look at the question of *when* to switch into exploration mode, and for how long.

119, TITLE: GreaseLM: Graph REASoning Enhanced Language Models
https://openreview.net/forum?id=41e9o6cQPj
AUTHORS: Xikun Zhang, Antoine Bosselut, Michihiro Yasunaga, Hongyu Ren, Percy Liang, Christopher D Manning, Jure Leskovec
HIGHLIGHT: We propose GreaseLM, a new model that fuses encoded representations from pretrained LMs and GNNs over multiple layers of modality interaction operations, allowing both modalities to bidirectionally inform the representation of the other.

120, TITLE: POETREE: Interpretable Policy Learning with Adaptive Decision Trees
https://openreview.net/forum?id=AJsI-ymaKn_
AUTHORS: Aliz?e Pace, Alex Chan, Mihaela van der Schaar
HIGHLIGHT: Policy Extraction through decision Trees (POETREE) is a novel framework for interpretable policy learning, compatible with fully-offline and partially-observable clinical decision environments.

121, TITLE: Convergent Boosted Smoothing for Modeling GraphData with Tabular Node Features

https://openreview.net/forum?id=nHpzE7DqAnG
AUTHORS: Jiuhai Chen, Jonas Mueller, Vassilis N. Ioannidis, Soji Adeshina, Yangkun Wang, Tom Goldstein, David Wipf
HIGHLIGHT: We develop a convergent method for combining boosting and graph propagation layers.

122, TITLE: Anomaly Transformer: Time Series Anomaly Detection with Association Discrepancy
https://openreview.net/forum?id=LzQQ89U1qm_
AUTHORS: Jiehui Xu, Haixu Wu, Jianmin Wang, Mingsheng Long
HIGHLIGHT: This paper detects time series anomalies from a new association-based dimension. We find an inherently normal-abnormal distinguishable evidence as Association Discrepancy. Co-designed with this evidence, our model achieves the SOTA on six benchmarks.

123, TITLE: Implicit Bias of Projected Subgradient Method Gives Provable Robust Recovery of Subspaces of Unknown Codimension
https://openreview.net/forum?id=vA7doMdgi75
AUTHORS: Paris Giampouras, Benjamin David Haeffele, Rene Vidal
HIGHLIGHT: We study the robust subspace recovery problem when subspace codimension is unknown.

124, TITLE: How Do Vision Transformers Work?
https://openreview.net/forum?id=D78Go4hVcxO
AUTHORS: Namuk Park, Songkuk Kim
HIGHLIGHT: We show that (1) multi-head self-attentions (MSAs) for computer vision flatten the loss landscapes, (2) MSAs are low-pass filters as opposed to Convs, and (3) MSAs at the end of a stage significantly improve the accuracy.

125, TITLE: Dynamics-Aware Comparison of Learned Reward Functions
https://openreview.net/forum?id=CALFyKVs87
AUTHORS: Blake Wulfe, Logan Michael Ellis, Jean Mercat, Rowan Thomas McAllister, Adrien Gaidon
HIGHLIGHT: We propose a method for quantifying the similarity of learned reward functions without performing policy learning and evaluation.

126, TITLE: What Happens after SGD Reaches Zero Loss? --A Mathematical Framework
https://openreview.net/forum?id=siCt4xZn5Ve
AUTHORS: Zhiyuan Li, Tianhao Wang, Sanjeev Arora
HIGHLIGHT: We propose a mathematical framework to study the implicit bias of SGD after reaching zero loss, based on which we prove label noise can help SGD escape the kernel regime and achieve optimal sample complexity for overparametrized linear model.

127, TITLE: R5: Rule Discovery with Reinforced and Recurrent Relational Reasoning
https://openreview.net/forum?id=2eXhNpHeW6E
AUTHORS: Shengyao Lu, Bang Liu, Keith G Mills, SHANGLING JUI, Di Niu
HIGHLIGHT: In this paper, we propose R5, a relational reasoning framework based on reinforcement learning that reasons over relational graph data and explicitly mines underlying compositional logical rules from observations.

128, TITLE: RelaxLoss: Defending Membership Inference Attacks without Losing Utility
https://openreview.net/forum?id=FEDfGWVZYIn
AUTHORS: Dingfan Chen, Ning Yu, Mario Fritz
HIGHLIGHT: We propose a novel training scheme that is highly effective in protecting against membership inference attacks while preserving the utility of target models.

129, TITLE: Independent SE(3)-Equivariant Models for End-to-End Rigid Protein Docking
https://openreview.net/forum?id=GQjaI9mLet
AUTHORS: Octavian-Eugen Ganea, Xinyuan Huang, Charlotte Bunne, Yatao Bian, Regina Barzilay, Tommi S. Jaakkola, Andreas Krause
HIGHLIGHT: We perform rigid protein docking using a novel independent SE(3)-equivariant message passing mechanism that guarantees the same resulting protein complex independent of the initial placement of the two 3D structures.

130, TITLE: Path Auxiliary Proposal for MCMC in Discrete Space
https://openreview.net/forum?id=JSR-YDImK95
AUTHORS: Haoran Sun, Hanjun Dai, Wei Xia, Arun Ramamurthy
HIGHLIGHT: In this work, we present a path auxiliary algorithm that uses a composition of local moves to efficiently explore large neighborhoods.

131, TITLE: Geometric and Physical Quantities improve E(3) Equivariant Message Passing
https://openreview.net/forum?id=_xwr8gOBeV1
AUTHORS: Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J Bekkers, Max Welling
HIGHLIGHT: We generalise equivariant graph networks such that node and edge updates are able to leverage covariant information.

132, TITLE: Learning Strides in Convolutional Neural Networks
https://openreview.net/forum?id=M752z9FKJP
AUTHORS: Rachid Riad, Olivier Teboul, David Grangier, Neil Zeghidour
HIGHLIGHT: We introduce DiffStride, the first downsampling layer with learnable strides for convolutional neural networks.

133, TITLE: VAE Approximation Error: ELBO and Exponential Families
https://openreview.net/forum?id=OIs3SxU5Ynl
AUTHORS: Alexander Shekhovtsov, Dmitrij Schlesinger, Boris Flach
HIGHLIGHT: VAEs have an inductive bias towards RBMs and generalized linear models

134, TITLE: Equivariant Subgraph Aggregation Networks
https://openreview.net/forum?id=dFbKQaRk15w
AUTHORS: Beatrice Bevilacqua, Fabrizio Frasca, Derek Lim, Balasubramaniam Srinivasan, Chen Cai, Gopinath Balamurugan, Michael M. Bronstein, Haggai Maron
HIGHLIGHT: We present a provably expressive graph learning framework based on representing graphs as multisets of subgraphs and processing them with an equivariant architecture.

135, TITLE: Source-Free Adaptation to Measurement Shift via Bottom-Up Feature Restoration
https://openreview.net/forum?id=1JDiK_TbV4S
AUTHORS: Cian Eastwood, Ian Mason, Chris Williams, Bernhard Sch?lkopf
HIGHLIGHT: We identify a type of domain shift which can be resolved by restoring the *same* features and address it in the source-free setting by using softly-binned histograms to cheaply and flexibly align the marginal feature distributions.

136, TITLE: On Lottery Tickets and Minimal Task Representations in Deep Reinforcement Learning
https://openreview.net/forum?id=Fl3Mg_MZR-
AUTHORS: Marc Vischer, Robert Tjarko Lange, Henning Sprekeler
HIGHLIGHT: We investigate the mechanisms underlying the lottery ticket effect in Deep RL and show that the derived mask extracts minimal task representations.

137, TITLE: Evaluation Metrics for Graph Generative Models: Problems, Pitfalls, and Practical Solutions
https://openreview.net/forum?id=tBtoZYKd9n
AUTHORS: Leslie O'Bray, Max Horn, Bastian Rieck, Karsten Borgwardt
HIGHLIGHT: We investigate the potential pitfalls of using MMD to evaluate graph generative models and propose recommendations for the practitioner on how to mitigate those challenges.

138, TITLE: Linking Emergent and Natural Languages via Corpus Transfer
https://openreview.net/forum?id=49A1Y6tRhaq
AUTHORS: Shunyu Yao, Mo Yu, Yang Zhang, Karthik R Narasimhan, Joshua B. Tenenbaum, Chuang Gan
HIGHLIGHT: We find that pre-training on an emergent language corpus improves natural language tasks in a low resource setup, and propose a metric to predict such a transferability.

139, TITLE: Interpretable Unsupervised Diversity Denoising and Artefact Removal
https://openreview.net/forum?id=DfMqlB0PXjM
AUTHORS: Mangal Prakash, Mauricio Delbracio, Peyman Milanfar, Florian Jug
HIGHLIGHT: This work proposes a new architecture for unsupervised, interpretable and diverse image restoration while achieving state-of-the-art results on numerous commonly used benchmarks across multiple image domains.

140, TITLE: Properties from mechanisms: an equivariance perspective on identifiable representation learning
https://openreview.net/forum?id=g5ynW-jMq4M
AUTHORS: Kartik Ahuja, Jason Hartford, Yoshua Bengio
HIGHLIGHT: Representation learning is identifiable up to any equivariances of the (known) mechanisms that govern an environment's evolution.

141, TITLE: Assessing Generalization of SGD via Disagreement
https://openreview.net/forum?id=WvOGCEAQhxl
AUTHORS: Yiding Jiang, Vaishnavh Nagarajan, Christina Baek, J Zico Kolter
HIGHLIGHT: We provide a surprisingly simple technique to accurately estimate the test error of deep neural networks using unlabeled data and we prove that this works because SGD ensembles are naturally well-calibrated.

142, TITLE: Scalable Sampling for Nonsymmetric Determinantal Point Processes
https://openreview.net/forum?id=BB4e8Atc1eR
AUTHORS: Insu Han, Mike Gartrell, Jennifer Gillenwater, Elvis Dohmatob, amin karbasi
HIGHLIGHT: We propose the first scalable linear-time and sublinear-time sampling algorithms for nonsymmetric determinantal point processes.

143, TITLE: The MultiBERTs: BERT Reproductions for Robustness Analysis
https://openreview.net/forum?id=K0E_F0gFDgA
AUTHORS: Thibault Sellam, Steve Yadlowsky, Ian Tenney, Jason Wei, Naomi Saphra, Alexander D'Amour, Tal Linzen, Jasmijn Bastings, Iulia Raluca Turc, Jacob Eisenstein, Dipanjan Das, Ellie Pavlick
HIGHLIGHT: We introduce MultiBERTs, 25 BERT checkpoints trained with similar hyper-parameters but different random seeds, and the Multi-Bootstrap, a bootstrapping method for experimental settings that involve multiple models and limited test data.

144, TITLE: Multitask Prompted Training Enables Zero-Shot Task Generalization
https://openreview.net/forum?id=9Vrb9D0WI4
AUTHORS: Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, Alexander M Rush
HIGHLIGHT: Can zero-shot generalization instead be directly induced by explicit multitask learning? To test this question at scale, we develop a system for easily mapping general natural language tasks into a human-readable prompted form.

145, TITLE: Self-Supervision Enhanced Feature Selection with Correlated Gates
https://openreview.net/forum?id=oDFvtxzPOx
AUTHORS: Changhee Lee, Fergus Imrie, Mihaela van der Schaar
HIGHLIGHT: We propose a novel DL-based feature selection method using self-supervised learning and multivariate Bernoulli distribution to address common challenges in feature selection: a scarcity of labeled samples and significant correlations among features.

146, TITLE: DEPTS: Deep Expansion Learning for Periodic Time Series Forecasting
https://openreview.net/forum?id=AJAR-JgNw__
AUTHORS: Wei Fan, Shun Zheng, Xiaohan Yi, Wei Cao, Yanjie Fu, Jiang Bian, Tie-Yan Liu
HIGHLIGHT: In this paper, we introduce a deep expansion learning framework, DEPTS, for PTS forecasting.

147, TITLE: Ab-Initio Potential Energy Surfaces by Pairing GNNs with Neural Wave Functions
https://openreview.net/forum?id=apv504XsysP
AUTHORS: Nicholas Gao, Stephan G?nnemann
HIGHLIGHT: We introduce a PESNet, a new network architecture that solves the Schr?dinger equation for multiple geometries simultaneously.

148, TITLE: Continuous-Time Meta-Learning with Forward Mode Differentiation
https://openreview.net/forum?id=57PipS27Km
AUTHORS: Tristan Deleu, David Kanaa, Leo Feng, Giancarlo Kerg, Yoshua Bengio, Guillaume Lajoie, Pierre-Luc Bacon
HIGHLIGHT: COMLN is a new meta-learning algorithm, where adaptation follows a gradient flow. It enables learning the amount of adaptation using SGD. We devise a novel efficient algorithm to compute the meta-gradients of COMLN, based on forward-mode diff.

149, TITLE: CoBERL: Contrastive BERT for Reinforcement Learning
https://openreview.net/forum?id=sRZ3GhmegS
AUTHORS: Andrea Banino, Adria Puigdomenech Badia, Jacob C Walker, Tim Scholtes, Jovana Mitrovic, Charles Blundell
HIGHLIGHT: A new loss and an improved architecture to efficiently train attentional models in reinforcement learning.

150, TITLE: Label Encoding for Regression Networks

https://openreview.net/forum?id=8WawVDdKqlL
AUTHORS: Deval Shah, Zi Yu Xue, Tor Aamodt
HIGHLIGHT: We propose binary-encoded labels (BEL) which improves regression by generalizing the application of binary classification.

151, TITLE: Online Hyperparameter Meta-Learning with Hypergradient Distillation
https://openreview.net/forum?id=01AMRlen9wJ
AUTHORS: Hae Beom Lee, Hayeon Lee, JaeWoong Shin, Eunho Yang, Timothy Hospedales, Sung Ju Hwang
HIGHLIGHT: We propose a gradient-based hyperparameter optimization method based on the idea of knowledge distillation, which is fully online and applicable to high-dimensional hyperparameters.

152, TITLE: Message Passing Neural PDE Solvers
https://openreview.net/forum?id=vSix3HPYKSU
AUTHORS: Johannes Brandstetter, Daniel E. Worrall, Max Welling
HIGHLIGHT: This paper introduces a message passing neural PDE solver that replaces all heuristically designed components in numerical PDE solvers with backprop-optimized neural function approximators.

153, TITLE: Leveraging Automated Unit Tests for Unsupervised Code Translation
https://openreview.net/forum?id=cmt-6KtR4c4
AUTHORS: Baptiste Roziere, Jie Zhang, Francois Charton, Mark Harman, Gabriel Synnaeve, Guillaume Lample
HIGHLIGHT: We leverage automatically created multilingual unit tests to improve unsupervised machine translation methods for source code and substantially outperform the state-of-the-art on all the language pairs we consider.

154, TITLE: On the approximation properties of recurrent encoder-decoder architectures
https://openreview.net/forum?id=xDIvIqQ3DXD
AUTHORS: Zhong Li, Haotian Jiang, Qianxiao Li
HIGHLIGHT: Approximation properties of recurrent encoder-decoder architectures are given, where the formed temporal product structure further characterises temporal relationships able to be efficiently learned.

155, TITLE: Relational Multi-Task Learning: Modeling Relations between Data and Tasks
https://openreview.net/forum?id=8Py-W8lSUgy
AUTHORS: Kaidi Cao, Jiaxuan You, Jure Leskovec
HIGHLIGHT: We propose MetaLink to solve a variety of multi-task learning settings, by constructing a knowledge graph over data points and tasks.

156, TITLE: Learning-Augmented k-means Clustering
https://openreview.net/forum?id=X8cLTHexYyY
AUTHORS: Jon Ergun, Zhili Feng, Sandeep Silwal, David Woodruff, Samson Zhou
HIGHLIGHT: We study the k-means problem augmented with a learning-based predictor that gives noisy information about true labels.

157, TITLE: On the Uncomputability of Partition Functions in Energy-Based Sequence Models
https://openreview.net/forum?id=SsPCtEY6yCl
AUTHORS: Chu-Cheng Lin, Arya D. McCarthy
HIGHLIGHT: EBMs over sequences have several theoretical limitations as learnable probabilistic sequence models.

158, TITLE: When Vision Transformers Outperform ResNets without Pre-training or Strong Data Augmentations
https://openreview.net/forum?id=LtKcMgGOeLt
AUTHORS: Xiangning Chen, Cho-Jui Hsieh, Boqing Gong
HIGHLIGHT: Hence, this paper investigates ViTs and MLP-Mixers from the lens of loss geometry, intending to improve the models' data efficiency at training and generalization at inference.

159, TITLE: EE-Net: Exploitation-Exploration Neural Networks in Contextual Bandits
https://openreview.net/forum?id=X_ch3VrNSRg
AUTHORS: Yikun Ban, Yuchen Yan, Arindam Banerjee, Jingrui He
HIGHLIGHT: In this paper, we propose "EE-Net," a neural-based bandit approach with a novel exploration strategy.

160, TITLE: D-CODE: Discovering Closed-form ODEs from Observed Trajectories
https://openreview.net/forum?id=wENMvIsxNN
AUTHORS: Zhaozhi Qian, Krzysztof Kacprzyk, Mihaela van der Schaar

HIGHLIGHT: In this work, we propose the Discovery of Closed-form ODE framework (D-CODE), which advances symbolic regression beyond the paradigm of supervised learning.

161, TITLE: TAMP-S2GCNets: Coupling Time-Aware Multipersistence Knowledge Representation with Spatio-Supra Graph Convolutional Networks for Time-Series Forecasting
https://openreview.net/forum?id=wv6g8fWLX2q
AUTHORS: Yuzhou Chen, Ignacio Segovia-Dominguez, Baris Coskunuzer, Yulia Gel
HIGHLIGHT: We make the first step toward integrating two emerging directions, time-aware deep learning and multi-parameter persistence, allowing us to infer latent time-conditioned relations among entities in multivariate time series forecasting tasks.

162, TITLE: Contrastive Fine-grained Class Clustering via Generative Adversarial Networks
https://openreview.net/forum?id=XWODe7ZLn8f
AUTHORS: Yunji Kim, Jung-Woo Ha
HIGHLIGHT: We proposed a method for unsupervised fine-grained class clustering that leverages the information-theoretic regularization term based on contrastive loss.

163, TITLE: Policy improvement by planning with Gumbel
https://openreview.net/forum?id=bERaNdoegnO
AUTHORS: Ivo Danihelka, Arthur Guez, Julian Schrittwieser, David Silver
HIGHLIGHT: We redesign AlphaZero to keep improving even when training with a small number of simulations.

164, TITLE: Understanding Latent Correlation-Based Multiview Learning and Self-Supervision: An Identifiability Perspective
https://openreview.net/forum?id=5FUq05QRc5b
AUTHORS: Qi Lyu, Xiao Fu, Weiran Wang, Songtao Lu
HIGHLIGHT: This work aims to understand latent correlation maximization-based deep multiview learning from a latent component identification viewpoint.

165, TITLE: A General Analysis of Example-Selection for Stochastic Gradient Descent
https://openreview.net/forum?id=7gWSJrP3opB
AUTHORS: Yucheng Lu, Si Yi Meng, Christopher De Sa
HIGHLIGHT: In this paper, we develop a broad condition on the sequence of examples used by SGD that is sufficient to prove tight convergence rates in both strongly convex and non-convex settings.

166, TITLE: NASPY: Automated Extraction of Automated Machine Learning Models
https://openreview.net/forum?id=KhLK0sHMgXK
AUTHORS: Xiaoxuan Lou, Shangwei Guo, Jiwei Li, Yaoxin Wu, Tianwei Zhang
HIGHLIGHT: We present NASPY, an end-to-end adversarial framework to extract the networkarchitecture of deep learning models from Neural Architecture Search (NAS).

167, TITLE: Strength of Minibatch Noise in SGD
https://openreview.net/forum?id=uorVGbWV5sw
AUTHORS: Liu Ziyin, Kangqiao Liu, Takashi Mori, Masahito Ueda
HIGHLIGHT: We solve the strength and shape of the minibatch noise in SGD exactly.

168, TITLE: Towards Deployment-Efficient Reinforcement Learning: Lower Bound and Optimality
https://openreview.net/forum?id=ccWaPGl9Hq
AUTHORS: Jiawei Huang, Jinglin Chen, Li Zhao, Tao Qin, Nan Jiang, Tie-Yan Liu
HIGHLIGHT: We propose a formal theoretical formulation for depolyment-efficient reinforcement learning; establish lower bounds for deployment complexity and study near-optimal deployment-efficient algorithms in linear MDP setting.

169, TITLE: Tackling the Generative Learning Trilemma with Denoising Diffusion GANs
https://openreview.net/forum?id=JprM0p-q0Co
AUTHORS: Zhisheng Xiao, Karsten Kreis, Arash Vahdat
HIGHLIGHT: To reduce the number of sampling steps in diffusion models, we propose to model the denoising distribution with conditional GANs. We show our model tackles the generative learning trilemma & achieves high sample quality, diversity & fast sampling.

170, TITLE: Tighter Sparse Approximation Bounds for ReLU Neural Networks

https://openreview.net/forum?id=LBvk4QWIUpm
AUTHORS: Carles Domingo-Enrich, Youssef Mroueh
HIGHLIGHT: We show conditions under which a function can be represented by an infinite-width neural network on a bounded set, and refine sparse neural network approximation bounds.

171, TITLE: Universal Approximation Under Constraints is Possible with Transformers
https://openreview.net/forum?id=JGO8CvG5S9
AUTHORS: Anastasis Kratsios, Behnoosh Zamanlooy, Tianlin Liu, Ivan Dokmanic
HIGHLIGHT: We provide the first universal approximation theorem with exact non-convex constraint satisfaction, and we introduce probabilistic transformer networks to do so.

172, TITLE: SOSP: Efficiently Capturing Global Correlations by Second-Order Structured Pruning
https://openreview.net/forum?id=t5EmXZ3ZLR
AUTHORS: Manuel Nonnenmacher, Thomas Pfeil, Ingo Steinwart, David Reeb
HIGHLIGHT: We introduce a second-order structured pruning method which efficiently captures global correlations among structures of deep neural networks.

173, TITLE: Representation Learning for Online and Offline RL in Low-rank MDPs
https://openreview.net/forum?id=J4iSIR9fhY0
AUTHORS: Masatoshi Uehara, Xuezhou Zhang, Wen Sun
HIGHLIGHT: We study representation learning in low-rank MDP in both online setting and offline setting, and propose statistically and computationally efficient algorithms.

174, TITLE: ViTGAN: Training GANs with Vision Transformers
https://openreview.net/forum?id=dwg5rXg1WS_
AUTHORS: Kwonjoon Lee, Huiwen Chang, Lu Jiang, Han Zhang, Zhuowen Tu, Ce Liu
HIGHLIGHT: Recently, Vision Transformers (ViTs) have shown competitive performance on image recognition while requiring less vision-specific inductive biases. In this paper, we investigate if such performance can be extended to image generation.

175, TITLE: Deconstructing the Inductive Biases of Hamiltonian Neural Networks
https://openreview.net/forum?id=EDeVYpT42oS
AUTHORS: Nate Gruver, Marc Anton Finzi, Samuel Don Stanton, Andrew Gordon Wilson
HIGHLIGHT: In this paper, we examine the inductive biases that make physics-inspired models successful in practice.

176, TITLE: Transition to Linearity of Wide Neural Networks is an Emerging Property of Assembling Weak Models
https://openreview.net/forum?id=CyKHoKyvgnp
AUTHORS: Chaoyue Liu, Libin Zhu, Misha Belkin
HIGHLIGHT: Transition to linearity of wide neural networks is an emerging property of assembling weak models corresponding to individual neurons

177, TITLE: Churn Reduction via Distillation
https://openreview.net/forum?id=HbtFCX2PLq0
AUTHORS: Heinrich Jiang, Harikrishna Narasimhan, Dara Bahri, Andrew Cotter, Afshin Rostamizadeh
HIGHLIGHT: We show distillation is a principled and practical solution to churn reduction.

178, TITLE: EntQA: Entity Linking as Question Answering
https://openreview.net/forum?id=US2rTP5nm_
AUTHORS: Wenzheng Zhang, Wenyue Hua, Karl Stratos
HIGHLIGHT: We frame entity linking as inverse open-domain question answering and solve the dilemma of having to predict mentions before entities.

179, TITLE: On the Connection between Local Attention and Dynamic Depth-wise Convolution
https://openreview.net/forum?id=L3_SsSNMmy
AUTHORS: Qi Han, Zejia Fan, Qi Dai, Lei Sun, Ming-Ming Cheng, Jiaying Liu, Jingdong Wang
HIGHLIGHT: We study the connection between local attention and dynamic depth-wise convolution in terms of sparse connectivity, weight sharing, and dynamic weight

180, TITLE: On the Optimal Memorization Power of ReLU Neural Networks
https://openreview.net/forum?id=MkTPtnjeYTV
AUTHORS: Gal Vardi, Gilad Yehudai, Ohad Shamir

HIGHLIGHT: We show that ReLU neural networks can memorize N samples using \sqrt{N} parameters, and prove that up to logarithmic terms this is the optimal solution.

181, TITLE: Neural Spectral Marked Point Processes
https://openreview.net/forum?id=0rcbOaoBXbg
AUTHORS: Shixiang Zhu, Haoyun Wang, Zheng Dong, Xiuyuan Cheng, Yao Xie
HIGHLIGHT: In this paper, we introduce a novel and general neural network-based non-stationary influence kernel with high expressiveness for handling complex discrete events data while providing theoretical performance guarantees.

182, TITLE: UniFormer: Unified Transformer for Efficient Spatial-Temporal Representation Learning
https://openreview.net/forum?id=nBU_u6DLvoK
AUTHORS: Kunchang Li, Yali Wang, Gao Peng, Guanglu Song, Yu Liu, Hongsheng Li, Yu Qiao
HIGHLIGHT: We introduce a novel Unified transFormer (UniFormer) which seamlessly integrates merits of 3D convolution and spatial-temporal self-attention in a concise transformer format, and achieves new state-of-the-art performances on Something-Something.

183, TITLE: Joint Shapley values: a measure of joint feature importance
https://openreview.net/forum?id=vcUmUvQCloe
AUTHORS: Chris Harris, Richard Pymar, Colin Rowat
HIGHLIGHT: We present a direct extension of Shapley's value to sets of features, thus extending the Shapley value's intuition: a set of feature's average effect on a model's prediction

184, TITLE: Understanding approximate and unrolled dictionary learning for pattern recovery
https://openreview.net/forum?id=rI0LYgGeYaw
AUTHORS: Beno?t Mal?zieux, Thomas Moreau, Matthieu Kowalski
HIGHLIGHT: This work studies an approximate formulation of dictionary learning based on unrolling and compares it to alternating minimization to find the best trade-off between speed and precision.

185, TITLE: Efficient Learning of Safe Driving Policy via Human-AI Copilot Optimization
https://openreview.net/forum?id=0cgU-BZp2ky
AUTHORS: Quanyi Li, Zhenghao Peng, Bolei Zhou
HIGHLIGHT: In this work, we develop a novel human-in-the-loop learning method called Human-AI Copilot Optimization (HACO).

186, TITLE: Evidential Turing Processes
https://openreview.net/forum?id=84NMXTHYe-
AUTHORS: Melih Kandemir, Abdullah Akg?l, Manuel Haussmann, Gozde Unal
HIGHLIGHT: An original extension of evidential deep learning with neural processes and neural Turing machines makes it possible to attain both in-domain calibration and out-of-domain detection in a single model.

187, TITLE: Trivial or Impossible --- dichotomous data difficulty masks model differences (on ImageNet and beyond)
https://openreview.net/forum?id=C_vsGwEIjAr
AUTHORS: Kristof Meding, Luca M. Schulze Buschoff, Robert Geirhos, Felix A. Wichmann
HIGHLIGHT: All CNNs make similar decisions since data difficulty is dichotomous.

188, TITLE: Associated Learning: an Alternative to End-to-End Backpropagation that Works on CNN, RNN, and Transformer
https://openreview.net/forum?id=4N-17dske79
AUTHORS: Dennis Y.H. Wu, Dinan Lin, Vincent Chen, Hung-Hsuan Chen
HIGHLIGHT: This paper studies Associate Learning, an alternative methodology to the end-to-end backpropagation

189, TITLE: How Attentive are Graph Attention Networks?
https://openreview.net/forum?id=F72ximsx7C1
AUTHORS: Shaked Brody, Uri Alon, Eran Yahav
HIGHLIGHT: We identify that Graph Attention Networks (GAT) compute a very weak form of attention. We show its empirical implications and propose a fix.

190, TITLE: Explaining Point Processes by Learning Interpretable Temporal Logic Rules
https://openreview.net/forum?id=P07dq7iSAGr
AUTHORS: Shuang Li, Mingquan Feng, Lu Wang, Abdelmajid Essofi, Yufeng Cao, Junchi Yan, Le Song

HIGHLIGHT: We propose a principled method to learn a set of human-readable logic rules to explain temporal point processes.

191, TITLE: Multi-Task Processes
https://openreview.net/forum?id=9otKVlgrpZG
AUTHORS: Donggyun Kim, Seongwoong Cho, Wonkwang Lee, Seunghoon Hong
HIGHLIGHT: We propose a new family of stochastic processes that can infer multiple heterogeneous functions jointly given a few incomplete observations (i.e., some functions may not be observed at each input).

192, TITLE: Lipschitz-constrained Unsupervised Skill Discovery
https://openreview.net/forum?id=BGvt0ghNgA
AUTHORS: Seohong Park, Jongwook Choi, Jaekyeom Kim, Honglak Lee, Gunhee Kim
HIGHLIGHT: We propose Lipschitz-constrained Skill Discovery (LSD), which encourages the agent to discover more dynamic and diverse skills without external rewards and additional prior knowledge, enabling zero-shot control on goal-reaching downstream tasks.

193, TITLE: NodePiece: Compositional and Parameter-Efficient Representations of Large Knowledge Graphs
https://openreview.net/forum?id=xMJWUKJnFSw
AUTHORS: Mikhail Galkin, Etienne Denis, Jiapeng Wu, William L. Hamilton
HIGHLIGHT: Node hashing in graphs for 10-100x embedding size reduction without significant performance losses on many tasks and inductive out of the box.

194, TITLE: Wish you were here: Hindsight Goal Selection for long-horizon dexterous manipulation
https://openreview.net/forum?id=FKp8-pIRo3y
AUTHORS: Todor Davchev, Oleg Olegovich Sushkov, Jean-Baptiste Regli, Stefan Schaal, Yusuf Aytar, Markus Wulfmeier, Jon Scholz
HIGHLIGHT: In this work, we extend hindsight relabelling mechanisms to guide exploration along task-specific distributions implied by a small set of successful demonstrations.

195, TITLE: Learning Efficient Online 3D Bin Packing on Packing Configuration Trees
https://openreview.net/forum?id=bfuGjlCwAq
AUTHORS: Hang Zhao, Yang Yu, Kai Xu
HIGHLIGHT: We propose to enhance the practical applicability of online 3D-BPP via learning on a hierarchical packing configuration tree which makes the DRL model easy to deal with practical constraints and well-performing even with continuous solution space.

196, TITLE: FedBABU: Toward Enhanced Representation for Federated Image Classification
https://openreview.net/forum?id=HuaYQfggn5u
AUTHORS: Jaehoon Oh, SangMook Kim, Se-Young Yun
HIGHLIGHT: We propose a novel algorithm, FedBABU, which updates and aggregates only the body during federated training for enhanced representation.

197, TITLE: Hybrid Random Features
https://openreview.net/forum?id=EMigfE6ZeS
AUTHORS: Krzysztof Marcin Choromanski, Han Lin, Haoxian Chen, Arijit Sehanobish, Yuanzhe Ma, Deepali Jain, Jake Varley, Andy Zeng, Michael S Ryoo, Valerii Likhosherstov, Dmitry Kalashnikov, Vikas Sindhwani, Adrian Weller
HIGHLIGHT: We propose a new class of random feature methods for softmax and Gaussian kernel estimation that are adaptable to provide particularly accurate approximation in the desired regions of interest.

198, TITLE: Should We Be Pre-training? An Argument for End-task Aware Training as an Alternative
https://openreview.net/forum?id=2bO2x8NAIMB
AUTHORS: Lucio M. Dery, Paul Michel, Ameet Talwalkar, Graham Neubig
HIGHLIGHT: When we know the end-task objective in advance, instead of pre-training on auxiliary objectives and then fine-tuning, we advocate for it to be introduced directly in training with auxiliary objectives

199, TITLE: Constructing a Good Behavior Basis for Transfer using Generalized Policy Updates
https://openreview.net/forum?id=7IWGzQ6gZ1D
AUTHORS: Safa Alver, Doina Precup
HIGHLIGHT: Specifically, we consider the framework of generalized policy evaluation and improvement, in which the rewards for all tasks of interest are assumed to be expressible as a linear combination of a fixed set of features.

200, TITLE: Adversarial Unlearning of Backdoors via Implicit Hypergradient
https://openreview.net/forum?id=MeeQkFYVbzW
AUTHORS: Yi Zeng, Si Chen, Won Park, Zhuoqing Mao, Ming Jin, Ruoxi Jia
HIGHLIGHT: A minimax formulation of backdoor removal and an implicit gradient-based solver surpasses the state-of-art methods' best results in higher efficacy, efficiency, robustness to variations in triggers, settings, poison ratio, and clean data size.

201, TITLE: Predicting Physics in Mesh-reduced Space with Temporal Attention
https://openreview.net/forum?id=XctLdNfCmP
AUTHORS: XU HAN, Han Gao, Tobias Pfaff, Jian-Xun Wang, Liping Liu
HIGHLIGHT: We use a GNN to locally summarize features and create coarsened, compact mesh representation of the system state, onto which we apply a transformer-style temporal attention module for physics prediction.

202, TITLE: Learning Synthetic Environments and Reward Networks for Reinforcement Learning
https://openreview.net/forum?id=C1_esHN6AVn
AUTHORS: Fabio Ferreira, Thomas Nierhoff, Andreas S?linger, Frank Hutter
HIGHLIGHT: We propose an evolution-based approach to meta-learn synthetic neural environments and reward neural networks for reinforcement learning.

203, TITLE: Controlling the Complexity and Lipschitz Constant improves Polynomial Nets
https://openreview.net/forum?id=dQ7Cy_ndl1s
AUTHORS: Zhenyu Zhu, Fabian Latorre, Grigorios Chrysos, Volkan Cevher
HIGHLIGHT: We provide sample complexity results and bounds on the Lipschitz constant of polynomial networks, which we use to construct a regularization scheme that improves the robustness against adversarial noise.

204, TITLE: Maximizing Ensemble Diversity in Deep Reinforcement Learning
https://openreview.net/forum?id=hjd-kcpDpf2
AUTHORS: Hassam Sheikh, Mariano Phielipp, Ladislau Boloni
HIGHLIGHT: Maximizing diversity in neural network improves performance ensemble based reinforcement learning

205, TITLE: Differentiable Prompt Makes Pre-trained Language Models Better Few-shot Learners
https://openreview.net/forum?id=ek9a0qIafW
AUTHORS: Ningyu Zhang, Luoqiu Li, Xiang Chen, Shumin Deng, Zhen Bi, Chuanqi Tan, Fei Huang, Huajun Chen
HIGHLIGHT: A differentiable prompt learning method for few-shot NLP with optimized prompt templates as well as labels.

206, TITLE: OntoProtein: Protein Pretraining With Gene Ontology Embedding
https://openreview.net/forum?id=yfe1VMYAXa4
AUTHORS: Ningyu Zhang, Zhen Bi, Xiaozhuan Liang, Siyuan Cheng, Haosen Hong, Shumin Deng, Qiang Zhang, Jiazhang Lian, Huajun Chen
HIGHLIGHT: A general framework to integrate knowledge graph (gene ontology) into protein pre-training.

207, TITLE: Likelihood Training of Schr?dinger Bridge using Forward-Backward SDEs Theory
https://openreview.net/forum?id=nioAdKCEdXB
AUTHORS: Tianrong Chen, Guan-Horng Liu, Evangelos Theodorou
HIGHLIGHT: We present a new computational framework, grounded on Forward-Backward SDEs theory, for the log-likelihood training of Schr?dinger Bridge and provide theoretical connections to score-baesd generative models.

208, TITLE: Is High Variance Unavoidable in RL? A Case Study in Continuous Control
https://openreview.net/forum?id=9xhgmsNVHu
AUTHORS: Johan Bjorck, Carla P Gomes, Kilian Q Weinberger
HIGHLIGHT: we study sources of variance in RL and propose methods to decrease it.

209, TITLE: Geometric Transformers for Protein Interface Contact Prediction
https://openreview.net/forum?id=CS4463zx6Hi
AUTHORS: Alex Morehead, Chen Chen, Jianlin Cheng
HIGHLIGHT: We introduce a geometry-evolving graph transformer for 3D protein structures and employ it to achieve state-of-the-art precision for predicting inter-protein residue-residue contacts in challenging protein complex targets.

210, TITLE: Few-shot Learning via Dirichlet Tessellation Ensemble
https://openreview.net/forum?id=6kCiVaoQdx9

AUTHORS: Chunwei Ma, Ziyun Huang, Mingchen Gao, Jinhui Xu
HIGHLIGHT: We developed a novel geometric framework that greatly improves few-shot classification, based on Cluster-induced Voronoi Diagram (CIVD).

211, TITLE: Task Affinity with Maximum Bipartite Matching in Few-Shot Learning
https://openreview.net/forum?id=u2GZOiUTbt
AUTHORS: Cat Phuoc Le, Juncheng Dong, Mohammadreza Soltani, Vahid Tarokh
HIGHLIGHT: Task affinity and its application in few-shot learning

212, TITLE: BAM: Bayes with Adaptive Memory
https://openreview.net/forum?id=NdOoQnYPj_
AUTHORS: Josue Nassar, Jennifer Rogers Brennan, Ben Evans, Kendall Lowrey
HIGHLIGHT: We augment Bayes with memory to generalize many frameworks and overcome limitations of traditional methods in non-stationary settings

213, TITLE: Adversarially Robust Conformal Prediction
https://openreview.net/forum?id=9L1BsI4wP1H
AUTHORS: Asaf Gendler, Tsui-Wei Weng, Luca Daniel, Yaniv Romano
HIGHLIGHT: Multi-class calibration procedure that is provably robust to adversarial attacks

214, TITLE: Creating Training Sets via Weak Indirect Supervision
https://openreview.net/forum?id=m8uJvVgwRci
AUTHORS: Jieyu Zhang, Bohan Wang, Xiangchen Song, Yujing Wang, Yaming Yang, Jing Bai, Alexander Ratner
HIGHLIGHT: In this work, we present a new weak supervision paradigm which automatically creates training sets for training a machine learning model given unlabeled dataset and indirect supervision sources.

215, TITLE: Learning Discrete Structured Variational Auto-Encoder using Natural Evolution Strategies
https://openreview.net/forum?id=JJCjv4dAbyL
AUTHORS: Alon Berliner, Guy Rotman, Yossi Adi, Roi Reichart, Tamir Hazan
HIGHLIGHT: In this work, we use Natural Evolution Strategies (NES), a class of gradient-free black-box optimization algorithms, to learn discrete structured VAEs.

216, TITLE: Pareto Policy Pool for Model-based Offline Reinforcement Learning
https://openreview.net/forum?id=OqcZu8JIIzS
AUTHORS: Yijun Yang, Jing Jiang, Tianyi Zhou, Jie Ma, Yuhui Shi
HIGHLIGHT: We propose a model-based offline RL method that builds a diverse set of optimal policies on Pareto front providing different levels of model return-uncertainty trade-off and it significantly outperforms single-policy methods.

217, TITLE: Do Not Escape From the Manifold: Discovering the Local Coordinates on the Latent Space of GANs
https://openreview.net/forum?id=aTzMi4yV_RO
AUTHORS: Jaewoong Choi, Junho Lee, Changyeon Yoon, Jung Ho Park, Geonho Hwang, Myungjoo Kang
HIGHLIGHT: We propose a method for finding local-geometry-aware traversal directions on the intermediate latent space of Generative Adversarial Networks (GANs).

218, TITLE: VOS: Learning What You Don't Know by Virtual Outlier Synthesis
https://openreview.net/forum?id=TW7d65uYu5M
AUTHORS: Xuefeng Du, Zhaoning Wang, Mu Cai, Yixuan Li
HIGHLIGHT: In this paper, we present VOS, a novel framework for OOD detection by adaptively synthesizing virtual outliers that can meaningfully regularize the model's decision boundary during training.

219, TITLE: Learning Representation from Neural Fisher Kernel with Low-rank Approximation
https://openreview.net/forum?id=J1rhANsCY9
AUTHORS: Ruixiang ZHANG, Shuangfei Zhai, Etai Littwin, Joshua M. Susskind
HIGHLIGHT: In this paper, we study the representation of neural networks from the view of kernels.

220, TITLE: Taming Sparsely Activated Transformer with Stochastic Experts
https://openreview.net/forum?id=B72HXs80q4
AUTHORS: Simiao Zuo, Xiaodong Liu, Jian Jiao, Young Jin Kim, Hany Hassan, Ruofei Zhang, Jianfeng Gao, Tuo Zhao
HIGHLIGHT: We propose a new variant of MoE, Transformer with Stochastic Experts, that is more parameter efficient.

221, TITLE: Fortuitous Forgetting in Connectionist Networks
https://openreview.net/forum?id=ei3SY1_zYsE
AUTHORS: Hattie Zhou, Ankit Vani, Hugo Larochelle, Aaron Courville
HIGHLIGHT: We introduce "forget-and-relearn" as a training paradigm where forgetting removes undesirable information and relearning bolsters useful features towards better generalization and compositionality.

222, TITLE: Multiset-Equivariant Set Prediction with Approximate Implicit Differentiation
https://openreview.net/forum?id=5K7RRqZEjoS
AUTHORS: Yan Zhang, David W Zhang, Simon Lacoste-Julien, Gertjan J. Burghouts, Cees G. M. Snoek
HIGHLIGHT: We propose a better permutation-equivariance property for multisets and improve an existing set predictor that has this property with approximate implicit differentiation

223, TITLE: Learnability Lock: Authorized Learnability Control Through Adversarial Invertible Transformations
https://openreview.net/forum?id=6VpeS27viTq
AUTHORS: Weiqi Peng, Jinghui Chen
HIGHLIGHT: To tackle this issue, this paper introduces and investigates a new concept called "learnability lock" for securing the process of data authorization.

224, TITLE: Differentiable Expectation-Maximization for Set Representation Learning
https://openreview.net/forum?id=MXdFBmHT4C
AUTHORS: Minyoung Kim
HIGHLIGHT: We propose a novel set embedding function, a feed-forward network defined as the (differentiable) maximum-a-posterior estimate of the mixture, approximately attained by a few Expectation-Maximization steps.

225, TITLE: Neural graphical modelling in continuous-time: consistency guarantees and algorithms
https://openreview.net/forum?id=SsHBkfeRF9L
AUTHORS: Alexis Bellot, Kim Branson, Mihaela van der Schaar
HIGHLIGHT: We present algorithms and consistency guarantees for graphical modelling in dynamical systems.

226, TITLE: GATSBI: Generative Adversarial Training for Simulation-Based Inference
https://openreview.net/forum?id=kR1hC6j48Tp
AUTHORS: Poornima Ramesh, Jan-Matthis Lueckmann, Jan Boelts, ?lvaro Tejero-Cantero, David S. Greenberg, Pedro J. Goncalves, Jakob H. Macke
HIGHLIGHT: Using generative adversarial networks for simulation-based inference

227, TITLE: Backdoor Defense via Decoupling the Training Process
https://openreview.net/forum?id=TySnJ-0RdKI
AUTHORS: Kunzhe Huang, Yiming Li, Baoyuan Wu, Zhan Qin, Kui Ren
HIGHLIGHT: We reveal that the hidden backdoors are embedded in the feature space mostly due to the end-to-end supervised training paradigm, based on which we propose a simple yet effective decoupling-based training method for backdoor defense.

228, TITLE: CURVATURE-GUIDED DYNAMIC SCALE NETWORKS FOR MULTI-VIEW STEREO
https://openreview.net/forum?id=_Wzj0J2xs2D
AUTHORS: Khang Truong Giang, Soohwan Song, Sungho Jo
HIGHLIGHT: This paper proposes a dynamic scale feature network to address the matching ambiguity problem in Multi-view stereo (MVS) and then designs an efficient MVS network to predict the depth maps.

229, TITLE: Scattering Networks on the Sphere for Scalable and Rotationally Equivariant Spherical CNNs
https://openreview.net/forum?id=bjy5Zb2fo2
AUTHORS: Jason McEwen, Christopher Wallis, Augustine N. Mavor-Parker
HIGHLIGHT: Scaling rotationally equivariant spherical CNNs to high-resolution data through spherical scattering networks

230, TITLE: Promoting Saliency From Depth: Deep Unsupervised RGB-D Saliency Detection
https://openreview.net/forum?id=BZnnMbt0pW
AUTHORS: Wei Ji, Jingjing Li, Qi Bi, chuan guo, Jie Liu, Li Cheng
HIGHLIGHT: We propose the first deep unsupervised RGB-D saliency detection method, which achieves appealing performance and does not require any human efforts compared to fully-supervised learning.

231, TITLE: A Tale of Two Flows: Cooperative Learning of Langevin Flow and Normalizing Flow Toward Energy-Based Model
https://openreview.net/forum?id=31d5RLCUuXC
AUTHORS: Jianwen Xie, Yaxuan Zhu, Jun Li, Ping Li
HIGHLIGHT: Joint learning of a short-run MCMC generator and a normalizing flow in the context of energy-based model for image representation and generation.

232, TITLE: Task-Induced Representation Learning
https://openreview.net/forum?id=OzyXtIZAzFv
AUTHORS: Jun Yamada, Karl Pertsch, Anisha Gunjal, Joseph J Lim
HIGHLIGHT: We introduce task-induced representation learning, which leverages task information in offline data from prior tasks to learn representations of visually complex scenes that model only task-relevant aspects and enable efficient learning of new tasks.

233, TITLE: Ancestral protein sequence reconstruction using a tree-structured Ornstein-Uhlenbeck variational autoencoder
https://openreview.net/forum?id=FZoZ7a31GCW
AUTHORS: Lys Sanz Moreta, Ola R?nning, Ahmad Salim Al-Sibahi, Jotun Hein, Douglas Theobald, Thomas Hamelryck
HIGHLIGHT: Ancestral protein sequence reconstruction using a tree-structured Ornstein-Uhlenbeck variational autoencoder

234, TITLE: On Robust Prefix-Tuning for Text Classification
https://openreview.net/forum?id=eBCmOocUejf
AUTHORS: Zonghan Yang, Yang Liu
HIGHLIGHT: We propose a robust prefix-tuning framework that improves robustness of prefix-tuning against different types of attacks while preserving its efficiency and modularity with interpretation from the perspective of optimal control.

235, TITLE: Gradient Information Matters in Policy Optimization by Back-propagating through Model
https://openreview.net/forum?id=rzvOQrnclO0
AUTHORS: Chongchong Li, Yue Wang, Wei Chen, Yuting Liu, Zhi-Ming Ma, Tie-Yan Liu
HIGHLIGHT: Considering the gradient information in the model learning is crucial for the model-based policy optimization according to our theoritical results. Motivated by such conclusion, we design a novel DDPPO algorithm that can achieve the SOTA performance.

236, TITLE: A Class of Short-term Recurrence Anderson Mixing Methods and Their Applications
https://openreview.net/forum?id=_X90SIKbHa
AUTHORS: Fuchao Wei, Chenglong Bao, Yang Liu
HIGHLIGHT: We develop a novel class of short-term recurrence Anderson mixing methods and validate its effectiveness in several applications including training neural networks.

237, TITLE: Measuring the Interpretability of Unsupervised Representations via Quantized Reversed Probing
https://openreview.net/forum?id=HFPTzdwN39
AUTHORS: Iro Laina, Yuki M Asano, Andrea Vedaldi
HIGHLIGHT: We propose quantized reverse probing as a information-theoretic measure to assess the degree to which self-supervised visual representations align with human-interpretable concepts.

238, TITLE: GraphENS: Neighbor-Aware Ego Network Synthesis for Class-Imbalanced Node Classification
https://openreview.net/forum?id=MXEl7i-iru
AUTHORS: Joonhyung Park, Jaeyun Song, Eunho Yang
HIGHLIGHT: In this paper, we hypothesize that overfitting to the neighbor sets of minor class due to message passing is a major challenge for class-imbalanced node classification.

239, TITLE: Revisit Kernel Pruning with Lottery Regulated Grouped Convolutions
https://openreview.net/forum?id=LdEhiMG9WLO
AUTHORS: Shaochen Zhong, Guanqun Zhang, Ningjia Huang, Shuai Xu
HIGHLIGHT: A simple yet effective structured pruning framework based on kernel pruning, weights shifting, and grouped convolutions.

240, TITLE: Illiterate DALL$\cdot$E Learns to Compose
https://openreview.net/forum?id=h0OYV0We3oh
AUTHORS: Gautam Singh, Fei Deng, Sungjin Ahn
HIGHLIGHT: To learn compositional slot-based representation of an image and perform slot composition for zero-shot novel image generation.

241, TITLE: PoNet: Pooling Network for Efficient Token Mixing in Long Sequences
https://openreview.net/forum?id=9jInD9JjicF
AUTHORS: Chao-Hong Tan, Qian Chen, Wen Wang, Qinglin Zhang, Siqi Zheng, Zhen-Hua Ling
HIGHLIGHT: We propose a novel Pooling Network for token mixing with linear complexity, achieve competitive performance on the Long Range Arena benchmark, and 95.7% of the accuracy of BERT on the GLUE demonstrating its transferability.

242, TITLE: Trigger Hunting with a Topological Prior for Trojan Detection
https://openreview.net/forum?id=TXsjU8BaibT
AUTHORS: Xiaoling Hu, Xiao Lin, Michael Cogswell, Yi Yao, Susmit Jha, Chao Chen
HIGHLIGHT: To this end, we propose innovative priors such as \textbf{diversity} and \textbf{topological simplicity} to not only increase the chances of finding the appropriate triggers but also improve the quality of the found triggers.

243, TITLE: Distributionally Robust Fair Principal Components via Geodesic Descents
https://openreview.net/forum?id=9NVd-DMtThY
AUTHORS: Hieu Vu, Toan Tran, Man-Chung Yue, Viet Anh Nguyen
HIGHLIGHT: In this paper, we propose a distributionally robust optimization problem for principal component analysis which internalizes a fairness criterion in the objective function.

244, TITLE: Data-Driven Offline Optimization for Architecting Hardware Accelerators
https://openreview.net/forum?id=GsH-K1VIyy
AUTHORS: Aviral Kumar, Amir Yazdanbakhsh, Milad Hashemi, Kevin Swersky, Sergey Levine
HIGHLIGHT: In this paper, we develop such a data-driven offline optimization method for designing hardware accelerators, dubbed PRIME, that enjoys all of these properties.

245, TITLE: Predictive Modeling in the Presence of Nuisance-Induced Spurious Correlations
https://openreview.net/forum?id=12RoR2o32T
AUTHORS: Aahlad Manas Puli, Lily H Zhang, Eric Karl Oermann, Rajesh Ranganath
HIGHLIGHT: This paper build models robust to nuisance-induced spurious correlations by constructing a representation that distills out the influence of the nuisance variables, while also maximizing its information with the label.

246, TITLE: End-to-End Learning of Probabilistic Hierarchies on Graphs
https://openreview.net/forum?id=g2LCQwG7Of
AUTHORS: Daniel Z?gner, Bertrand Charpentier, Morgane Ayle, Sascha Geringer, Stephan G?nnemann
HIGHLIGHT: End-to-end gradient-based hierarchical clustering on graphs by exploiting Markov chain theory leads to state-of-the-art results.

247, TITLE: Regularized Autoencoders for Isometric Representation Learning
https://openreview.net/forum?id=mQxt8l7JL04
AUTHORS: Yonghyeon LEE, Sangwoong Yoon, MinJun Son, Frank C. Park
HIGHLIGHT: Regularized Autoencoders that simultaneously learn data manifold and a set of latent space coordinates that preserves the geometry of the learned manifold.

248, TITLE: LIGS: Learnable Intrinsic-Reward Generation Selection for Multi-Agent Learning
https://openreview.net/forum?id=CpTuR2ECuW
AUTHORS: David Henry Mguni, Taher Jafferjee, Jianhong Wang, Nicolas Perez-Nieves, Oliver Slumbers, Feifei Tong, Yang Li, Jiangcheng Zhu, Yaodong Yang, Jun Wang
HIGHLIGHT: In this paper, we introduce a new general framework for improving coordination and performance of multi-agent reinforcement learners (MARL).

249, TITLE: Learning to Generalize across Domains on Single Test Samples
https://openreview.net/forum?id=CIaQKbTBwtU
AUTHORS: Zehao Xiao, Xiantong Zhen, Ling Shao, Cees G. M. Snoek
HIGHLIGHT: We leverage a meta-learning paradigm to learn our model to acquire the ability of adaptation with single samples at training time so as to further adapt itself to each single test sample at test time.

250, TITLE: Unsupervised Learning of Full-Waveform Inversion: Connecting CNN and Partial Differential Equation in a Loop
https://openreview.net/forum?id=izvwgBic9q

AUTHORS: Peng Jin, Xitong Zhang, Yinpeng Chen, Sharon X Huang, Zicheng Liu, Youzuo Lin
HIGHLIGHT: We develop an unsupervised method to solve seismic full-waveform inversion in geophysics by integrating CNN and the governing partial differential equation.
We also introduce a new large-scale dataset $\textit{OpenFWI}$, to establish a more challenging benchmark for the community.

251, TITLE: On the Convergence of the Monte Carlo Exploring Starts Algorithm for Reinforcement Learning
https://openreview.net/forum?id=JzNB0eA2-M4
AUTHORS: Che Wang, Shuhan Yuan, Kai Shao, Keith W. Ross
HIGHLIGHT: We prove that the Monte Carlo Exploring Starts algorithm converges for optimal policy feed-forward MDPs.

252, TITLE: Deep Ensembling with No Overhead for either Training or Testing: The All-Round Blessings of Dynamic Sparsity
https://openreview.net/forum?id=RLtqs6pzj1-
AUTHORS: Shiwei Liu, Tianlong Chen, Zahra Atashgahi, Xiaohan Chen, Ghada Sokar, Elena Mocanu, Mykola Pechenizkiy, Zhangyang Wang, Decebal Constantin Mocanu
HIGHLIGHT: We propose an efficient ensemble learning framework FreeTickets via dynamic spasity, which is more efficient to train and inference than a single dense model, while matching the performance of the naive dense ensemble.

253, TITLE: Dive Deeper Into Integral Pose Regression
https://openreview.net/forum?id=vHVcB-ak3Si
AUTHORS: Kerui Gu, Linlin Yang, Angela Yao
HIGHLIGHT: We do a deep dive on the inference and back-propagation of integral pose regression to better understand the causes behind the performance and training differences.

254, TITLE: Retriever: Learning Content-Style Representation as a Token-Level Bipartite Graph
https://openreview.net/forum?id=AXWygMvuT6Q
AUTHORS: Dacheng Yin, Xuanchi Ren, Chong Luo, Yuwang Wang, Zhiwei Xiong, Wenjun Zeng
HIGHLIGHT: We propose a model-agnostic and unsupervised framework to learn a novel token-level bipartite graph representation of content and style from structured input.

255, TITLE: Do Users Benefit From Interpretable Vision? A User Study, Baseline, And Dataset
https://openreview.net/forum?id=v6s3HVjPerv
AUTHORS: Leon Sixt, Martin Schuessler, Oana-Iuliana Popescu, Philipp Wei?, Tim Landgraf
HIGHLIGHT: Do Users Benefit From Interpretable Vision? A User Study, Baseline, And Dataset

256, TITLE: Open-vocabulary Object Detection via Vision and Language Knowledge Distillation
https://openreview.net/forum?id=lL3lnMbR4WU
AUTHORS: Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, Yin Cui
HIGHLIGHT: We propose using knowledge distillation to train an object detector that can detect objects with arbitrary text inputs, outperforming its supervised counterparts on rare categories.

257, TITLE: Meta Learning Low Rank Covariance Factors for Energy Based Deterministic Uncertainty
https://openreview.net/forum?id=GQd7mXSPua
AUTHORS: Jeffrey Ryan Willette, Hae Beom Lee, Juho Lee, Sung Ju Hwang
HIGHLIGHT: We propose a novel meta learning algorithm which learns low rank covariance factors, and utilizes an energy-based inference to achieve a calibrated prediction.

258, TITLE: CKConv: Continuous Kernel Convolution For Sequential Data
https://openreview.net/forum?id=8FhxBtXSl0
AUTHORS: David W. Romero, Anna Kuzina, Erik J Bekkers, Jakub Mikolaj Tomczak, Mark Hoogendoorn
HIGHLIGHT: We provide a continuous parameterization to convolutional kernels, with which several advantages upon conventional (discrete) parameterizations are obtained.

259, TITLE: PriorGrad: Improving Conditional Denoising Diffusion Models with Data-Dependent Adaptive Prior
https://openreview.net/forum?id=_BNiN4IjC5
AUTHORS: Sang-gil Lee, Heeseung Kim, Chaehun Shin, Xu Tan, Chang Liu, Qi Meng, Tao Qin, Wei Chen, Sungroh Yoon, Tie-Yan Liu
HIGHLIGHT: We improve the efficiency of diffusion-based conditional generative models for audio by using data-dependent non-standard Gaussian as a prior.

260, TITLE: Patch-Fool: Are Vision Transformers Always Robust Against Adversarial Perturbations?
https://openreview.net/forum?id=28ib9tf6zhr
AUTHORS: Yonggan Fu, Shunyao Zhang, Shang Wu, Cheng Wan, Yingyan Lin
HIGHLIGHT: We propose the Patch-Fool attack to unveil a vulnerability perspective of ViTs.

261, TITLE: R4D: Utilizing Reference Objects for Long-Range Distance Estimation
https://openreview.net/forum?id=MQ2sAGunyBP
AUTHORS: Yingwei Li, Tiffany Chen, Maya Kabkab, Ruichi Yu, Longlong Jing, Yurong You, Hang Zhao
HIGHLIGHT: We then propose R4D, the first framework to accurately estimate the distance of long-range objects by using references with known distances in the scene.

262, TITLE: Eigencurve: Optimal Learning Rate Schedule for SGD on Quadratic Objectives with Skewed Hessian Spectrums
https://openreview.net/forum?id=rTAclwH46Tb
AUTHORS: Rui Pan, Haishan Ye, Tong Zhang
HIGHLIGHT: A learning rate schedule which achieves minimax optimal convergence rate (up to a constant) for SGD on quadratic objectives with skewed Hessian spectrums.

263, TITLE: Model-augmented Prioritized Experience Replay
https://openreview.net/forum?id=WuEiafqdy9H
AUTHORS: Youngmin Oh, Jinwoo Shin, Eunho Yang, Sung Ju Hwang
HIGHLIGHT: We propose a novel experience replay which employs additional auxiliary learnable features as well as TD-errors for prioritizing experiences

264, TITLE: Actor-Critic Policy Optimization in a Large-Scale Imperfect-Information Game
https://openreview.net/forum?id=DTXZqTNV5nW
AUTHORS: Haobo Fu, Weiming Liu, Shuang Wu, Yijia Wang, Tao Yang, Kai Li, Junliang Xing, Bin Li, Bo Ma, QIANG FU, Yang Wei
HIGHLIGHT: A new actor-critic algorithm for approximating a Nash Equilibrium in the large-scale imperfect-information game 1v1 Mahjong.

265, TITLE: Reducing the Communication Cost of Federated Learning through Multistage Optimization
https://openreview.net/forum?id=ZaVVVlcdaN
AUTHORS: Charlie Hou, Kiran Koshy Thekumparampil, Giulia Fanti, Sewoong Oh
HIGHLIGHT: In this paper, we propose a multistage optimization scheme that nearly matches the lower bound across all heterogeneity levels.

266, TITLE: Is Homophily a Necessity for Graph Neural Networks?
https://openreview.net/forum?id=ucASPPD9GKN
AUTHORS: Yao Ma, Xiaorui Liu, Neil Shah, Jiliang Tang
HIGHLIGHT: Our work carefully characterizes these conditions and provides supporting theoretical understanding and empirical observations.

267, TITLE: How Does SimSiam Avoid Collapse Without Negative Samples? Towards a Unified Understanding of Progress in SSL
https://openreview.net/forum?id=bwq6O4Cwdl
AUTHORS: Chaoning Zhang, Kang Zhang, Chenshuang Zhang, Trung X. Pham, Chang D. Yoo, In So Kweon
HIGHLIGHT: After refuting their claims, we introduce vector decomposition for analyzing the collapse based on the gradient analysis of l_2 normalized vector.

268, TITLE: Better Supervisory Signals by Observing Learning Paths
https://openreview.net/forum?id=Iog0djAdbHj
AUTHORS: Yi Ren, Shangmin Guo, Danica J. Sutherland
HIGHLIGHT: A study of how and why teachers in knowledge distillation end up with better supervisory signals than the original labels.

269, TITLE: Fooling Explanations in Text Classifiers
https://openreview.net/forum?id=j3krplz_4w6
AUTHORS: Adam Ivankay, Ivan Girardi, Chiara Marchiori, Pascal Frossard
HIGHLIGHT: Our work shows that explanation methods in text classifiers are susceptible to imperceptible perturbations that alter the explanation outcomes without changing the predictions of the classifiers.

270, TITLE: Certified Robustness for Deep Equilibrium Models via Interval Bound Propagation
https://openreview.net/forum?id=y1PXylgrXZ
AUTHORS: Colin Wei, J Zico Kolter
HIGHLIGHT: To develop certifiably robust deep equilibrium (DEQ) models, we propose the IBP-MonDEQ layer, a DEQ layer where interval bounds on the output can be obtained by solving an additional fixed-point equation inspired by interval bound propagation.

271, TITLE: LFPT5: A Unified Framework for Lifelong Few-shot Language Learning Based on Prompt Tuning of T5
https://openreview.net/forum?id=HCRVf71PMF
AUTHORS: Chengwei Qin, Shafiq Joty
HIGHLIGHT: We define a challenging yet practical problem as Lifelong Few-shot Language Learning and propose a unified framework for it based on prompt tuning of T5.

272, TITLE: Steerable Partial Differential Operators for Equivariant Neural Networks
https://openreview.net/forum?id=N9W24a4zU
AUTHORS: Erik Jenner, Maurice Weiler
HIGHLIGHT: We present a framework for equivariant partial differential operators, generalizing existing approaches and narrowing the gap between PDOs and convolutions.

273, TITLE: Divisive Feature Normalization Improves Image Recognition Performance in AlexNet
https://openreview.net/forum?id=aOX3a9q3RVV
AUTHORS: Michelle Miller, SueYeon Chung, Kenneth D. Miller
HIGHLIGHT: DIVISIVE FEATURE NORMALIZATION IMPROVES IMAGE RECOGNITION PERFORMANCE AND IN- CREASES MANIFOLD CAPACITY, SPARSITY, AND LOW-FREQUENCY REPRESENTATION IN DEEP NETS

274, TITLE: Actor-critic is implicitly biased towards high entropy optimal policies
https://openreview.net/forum?id=vEZyTBRPP6o
AUTHORS: Yuzheng Hu, Ziwei Ji, Matus Telgarsky
HIGHLIGHT: We show that actor-critic, without any explicit exploration or regularization, can obtain an ϵ-optimal high entropy policy in $\text{poly}(1/\epsilon)$ samples via a single trajectory without the usual uniform mixing assumptions.

275, TITLE: Top-N: Equivariant Set and Graph Generation without Exchangeability
https://openreview.net/forum?id=-Gk_IPJWvk
AUTHORS: Clement Vignac, Pascal Frossard
HIGHLIGHT: We propose the Top-N method for one-shot set and graph probabilistic decoders as a replacement for i.i.d. generation in the first layer.

276, TITLE: CrossMatch: Cross-Classifier Consistency Regularization for Open-Set Single Domain Generalization
https://openreview.net/forum?id=48RBsJwGkJf
AUTHORS: Ronghang Zhu, Sheng Li
HIGHLIGHT: In this paper, we propose a challenging and untouched problem: \textit{Open-Set Single Domain Generalization} (OS-SDG), where target domains include unseen categories out of source label space.

277, TITLE: Map Induction: Compositional spatial submap learning for efficient exploration in novel environments
https://openreview.net/forum?id=1NUsBU-7HAL
AUTHORS: Sugandha Sharma, Aidan Curtis, Marta Kryven, Joshua B. Tenenbaum, Ila R Fiete
HIGHLIGHT: Modelling Map Induction for efficient exploration in novel environments.

278, TITLE: Exposing the Implicit Energy Networks behind Masked Language Models via Metropolis--Hastings
https://openreview.net/forum?id=6PvWo1kEvlT
AUTHORS: Kartik Goyal, Chris Dyer, Taylor Berg-Kirkpatrick
HIGHLIGHT: We interpret masked language models for sequences as energy based models and propose a tractable scheme inspired by Metropolis--Hasting Monte Carlo to draw samples from these models.

279, TITLE: TAPEX: Table Pre-training via Learning a Neural SQL Executor
https://openreview.net/forum?id=O50443AsCP
AUTHORS: Qian Liu, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, Jian-Guang Lou
HIGHLIGHT: This work performs table pre-training by learning a neural SQL executor over a synthetic corpus, which is obtained by automatically synthesizing executable SQL queries and their execution results.

280, TITLE: From Stars to Subgraphs: Uplifting Any GNN with Local Structure Awareness
https://openreview.net/forum?id=Mspk_WYKoEH
AUTHORS: Lingxiao Zhao, Wei Jin, Leman Akoglu, Neil Shah
HIGHLIGHT: Our work stands between these two regimes: we introduce a general framework to uplift any MPNN to be more expressive, with limited scalability overhead and greatly improved practical performance.

281, TITLE: DISSECT: Disentangled Simultaneous Explanations via Concept Traversals
https://openreview.net/forum?id=qY79G8jGsep
AUTHORS: Asma Ghandeharioun, Been Kim, Chun-Liang Li, Brendan Jou, Brian Eoff, Rosalind Picard
HIGHLIGHT: We propose a novel counterfactual explainability method that simultaneously satisfies several desirable qualities where other methods fail by training a generator, a discriminator, and a concept disentangler using the classifier?s signal.

282, TITLE: Continuously Discovering Novel Strategies via Reward-Switching Policy Optimization
https://openreview.net/forum?id=hcQHRHKfN_
AUTHORS: Zihan Zhou, Wei Fu, Bingliang Zhang, Yi Wu
HIGHLIGHT: We propose Reward-Switching Policy Optimization (RSPO), a paradigm to discover diverse strategies in complex RL environments by iteratively finding novel policies that are both locally optimal and sufficiently different from existing ones.

283, TITLE: Sequence Approximation using Feedforward Spiking Neural Network for Spatiotemporal Learning: Theory and Optimization Methods
https://openreview.net/forum?id=bp-LJ4y_XC
AUTHORS: Xueyuan She, Saurabh Dash, Saibal Mukhopadhyay
HIGHLIGHT: A theoretical approache to study the approximation capability of feedforward spiking neural network and optimization methods for such network.

284, TITLE: Givens Coordinate Descent Methods for Rotation Matrix Learning in Trainable Embedding Indexes
https://openreview.net/forum?id=9-Rfew334N
AUTHORS: Yunjiang Jiang, Han Zhang, Yiming Qiu, Yun Xiao, Bo Long, Wen-Yun Yang
HIGHLIGHT: Learning orthonormal matrix in neural networks via Givens rotations

285, TITLE: Variational Predictive Routing with Nested Subjective Timescales
https://openreview.net/forum?id=JxFgJbZ-wft
AUTHORS: Alexey Zakharov, Qinghai Guo, Zafeirios Fountas
HIGHLIGHT: Variational inference hierarchical model that relies on a change detection mechanism to impose a nested temporal hierarchy on its latent structure.

286, TITLE: NASViT: Neural Architecture Search for Efficient Vision Transformers with Gradient Conflict aware Supernet Training
https://openreview.net/forum?id=Qaw16njk6L
AUTHORS: Chengyue Gong, Dilin Wang, Meng Li, Xinlei Chen, Zhicheng Yan, Yuandong Tian, qiang liu, Vikas Chandra
HIGHLIGHT: we identify one key issue of ViT supernet training that the supernet gradients and the sub-network gradients are likely to disagree with each other, and propose gradient conflict aware training.

287, TITLE: MCMC Should Mix: Learning Energy-Based Model with Flow-Based Backbone
https://openreview.net/forum?id=4C93Qvn-tz
AUTHORS: Erik Nijkamp, Ruiqi Gao, Pavel Sountsov, Srinivas Vasudevan, Bo Pang, Song-Chun Zhu, Ying Nian Wu
HIGHLIGHT: Learning energy-based models with mixing Markov chains.

288, TITLE: When, Why, and Which Pretrained GANs Are Useful?
https://openreview.net/forum?id=4Ycr8oeCoIh
AUTHORS: Timofey Grigoryev, Andrey Voynov, Artem Babenko
HIGHLIGHT: This work aims to dissect the process of GAN finetuning.

289, TITLE: Can an Image Classifier Suffice For Action Recognition?
https://openreview.net/forum?id=qhkFX-HLuHV
AUTHORS: Quanfu Fan, Chun-Fu Chen, Rameswar Panda
HIGHLIGHT: We propose the idea of super images to re-purpose an image classifer for action recognition.

290, TITLE: Graph-Enhanced Exploration for Goal-oriented Reinforcement Learning
https://openreview.net/forum?id=rlYiXFdSy70
AUTHORS: Jiarui Jin, Sijin Zhou, Weinan Zhang, Tong He, Yong Yu, Rasool Fakoor
HIGHLIGHT: In this paper, we propose G2RL, a new goal-oriented RL that leverages the state-transition graph for effective exploration and efficient training.

291, TITLE: Near-optimal Offline Reinforcement Learning with Linear Representation: Leveraging Variance Information with Pessimism
https://openreview.net/forum?id=KLaDXLAzzFT
AUTHORS: Ming Yin, Yaqi Duan, Mengdi Wang, Yu-Xiang Wang
HIGHLIGHT: Towards this goal, we study the statistical limits of offline reinforcement learning with linear model representations.

292, TITLE: Towards Model Agnostic Federated Learning Using Knowledge Distillation
https://openreview.net/forum?id=lQI_mZjvBxj
AUTHORS: Andrei Afonin, Sai Praneeth Karimireddy
HIGHLIGHT: We develop a rich yet tractable framework for analyzing distillation based federated learning algorithms, using which we draw some surprising insights.

293, TITLE: Modeling Label Space Interactions in Multi-label Classification using Box Embeddings
https://openreview.net/forum?id=tyTH9kOxcvh
AUTHORS: Dhruvesh Patel, Pavitra Dangati, Jay-Yoon Lee, Michael Boratko, Andrew McCallum
HIGHLIGHT: Improving the consistency for multi-label classification by modeling label space interactions using Box Embeddings.

294, TITLE: Anisotropic Random Feature Regression in High Dimensions
https://openreview.net/forum?id=JfaWawZ8BmX
AUTHORS: Gabriel Mel, Jeffrey Pennington
HIGHLIGHT: We derive exact asymptotic formulas for the total error, bias, and variance of random feature regression with anisotropic inputs and target weights, and identify a new type of singularity in sample-wise learning curves.

295, TITLE: Discrete Representations Strengthen Vision Transformer Robustness
https://openreview.net/forum?id=8hWs60AZcWk
AUTHORS: Chengzhi Mao, Lu Jiang, Mostafa Dehghani, Carl Vondrick, Rahul Sukthankar, Irfan Essa
HIGHLIGHT: We present a simple and effective architecture modification to ViT's input layer with discrete token representations.

296, TITLE: Confidence Adaptive Anytime Pixel-Level Recognition
https://openreview.net/forum?id=kNKFOXleuC
AUTHORS: Zhuang Liu, Hung-Ju Wang, Zhiqiu Xu, Trevor Darrell, Evan Shelhamer
HIGHLIGHT: First single-model anytime approach for pixel-level visual recognition; our model with redesigned exits and confidence adaptivity enables anytime inference, achieves the same level of final accuracy, and significantly reduces total computation.

297, TITLE: Automatic Loss Function Search for Predict-Then-Optimize Problems with Strong Ranking Property
https://openreview.net/forum?id=hSktDu-h94
AUTHORS: Boshi Wang, Jialin Yi, Hang Dong, Bo Qiao, Chuan Luo, Qingwei Lin
HIGHLIGHT: To properly connect the prediction loss with the optimization goal, in this paper we propose a total group preorder (TGP) loss and its differential version called approximated total group preorder (ATGP) loss for predict-then-optimize (PTO) problems with strong ranking property.

298, TITLE: Environment Predictive Coding for Visual Navigation
https://openreview.net/forum?id=DBiQQYWykyy
AUTHORS: Santhosh Kumar Ramakrishnan, Tushar Nagarajan, Ziad Al-Halah, Kristen Grauman
HIGHLIGHT: We introduce environment predicting coding, a self-supervised approach for learning environment-level representations for navigation-like tasks.

299, TITLE: Model Zoo: A Growing Brain That Learns Continually
https://openreview.net/forum?id=WfvgGBcgbE7
AUTHORS: Rahul Ramesh, Pratik Chaudhari

HIGHLIGHT: Continual learning methods can benefit by splitting the capacity of the learner and we leverage this in our method Model Zoo, which demonstrates large gains in accuracy on a variety of continual learning benchmarks.

300, TITLE: Online Adversarial Attacks
https://openreview.net/forum?id=bYGSzbCM_i
AUTHORS: Andjela Mladenovic, Joey Bose, Hugo berard, William L. Hamilton, Simon Lacoste-Julien, Pascal Vincent, Gauthier Gidel
HIGHLIGHT: We consider a new adversarial attack setting in which the data arrives as a stream and an adversary must pick top-k items to craft blackbox transfer attacks against an unkown target model.

301, TITLE: Distributional Reinforcement Learning with Monotonic Splines
https://openreview.net/forum?id=C8Ltz08PtBp
AUTHORS: Yudong Luo, Guiliang Liu, Haonan Duan, Oliver Schulte, Pascal Poupart
HIGHLIGHT: In this paper, we propose to learn smooth continuous quantile functions represented by monotonic rational-quadratic splines, which also naturally solve the quantile crossing problem.

302, TITLE: A Comparison of Variable Selection Methods for Blockwise Diagonal Designs
https://openreview.net/forum?id=nhN-fqxmNGx
AUTHORS: Tracy Ke, Longlin Wang
HIGHLIGHT: A theoretical comparison of the Hamming errors for 6 different variable selection methods

303, TITLE: RegionViT: Regional-to-Local Attention for Vision Transformers
https://openreview.net/forum?id=T__V3uLix7V
AUTHORS: Chun-Fu Chen, Rameswar Panda, Quanfu Fan
HIGHLIGHT: A new architecture for vision transformer

304, TITLE: What Do We Mean by Generalization in Federated Learning?
https://openreview.net/forum?id=VimqQq-i_Q
AUTHORS: Honglin Yuan, Warren Richard Morningstar, Lin Ning, Karan Singhal
HIGHLIGHT: We propose a framework for better measuring generalization and heterogeneity in federated learning, apply it for extensive empirical evaluation across six tasks, and make a series of recommendations for future FL works.

305, TITLE: From Intervention to Domain Transportation: A Novel Perspective to Optimize Recommendation
https://openreview.net/forum?id=jT1EwXu-4hj
AUTHORS: Da Xu, Yuting Ye, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, Kannan Achan
HIGHLIGHT: We propose and study a novel domain-transportation view for optimizing recommendation for information retrieval systems.

306, TITLE: Should I Run Offline Reinforcement Learning or Behavioral Cloning?
https://openreview.net/forum?id=AP1MKT37rJ
AUTHORS: Aviral Kumar, Joey Hong, Anikait Singh, Sergey Levine
HIGHLIGHT: Characterization of scenarios where offline reinforcement learning outperforms behavioral cloning

307, TITLE: Neural Stochastic Dual Dynamic Programming
https://openreview.net/forum?id=aisKPsMM3fg
AUTHORS: Hanjun Dai, Yuan Xue, Zia Syed, Dale Schuurmans, Bo Dai
HIGHLIGHT: We proposed neural-SDDP pushing the frontier of multi-stage stochastic optimization solver towards practical problem size.

308, TITLE: Incremental False Negative Detection for Contrastive Learning
https://openreview.net/forum?id=dDjSKKA5TP1
AUTHORS: Tsai-Shien Chen, Wei-Chih Hung, Hung-Yu Tseng, Shao-Yi Chien, Ming-Hsuan Yang
HIGHLIGHT: This paper explores the effect of false negative samples in self-supervised contrastive learning and introduce a framework to incrementally detect and explicitly remove the false negatives.

309, TITLE: Scaling the Depth of Vision Transformers via the Fourier Domain Analysis
https://openreview.net/forum?id=O476oWmiNNp
AUTHORS: Peihao Wang, Wenqing Zheng, Tianlong Chen, Zhangyang Wang
HIGHLIGHT: In this paper, we investigate the scalability issue with ViT via Fourier domain analysis and propose two practical solutions by scaling different frequency components of attention and feature maps.

310, TITLE: Deep AutoAugment
https://openreview.net/forum?id=St-53J9ZARf
AUTHORS: Yu Zheng, Zhi Zhang, Shen Yan, Mi Zhang
HIGHLIGHT: We propose Deep AutoAugment (DeepAA), a fully automated automated data augmentation methods that outperforms previous automated data augmentation methods.

311, TITLE: HyperDQN: A Randomized Exploration Method for Deep Reinforcement Learning
https://openreview.net/forum?id=X0nrKAXu7g-
AUTHORS: Ziniu Li, Yingru Li, Yushun Zhang, Tong Zhang, Zhi-Quan Luo
HIGHLIGHT: We design a practical randomized exploration method to address the sample efficiency issue in online reinforcement learning.

312, TITLE: Chaos is a Ladder: A New Understanding of Contrastive Learning
https://openreview.net/forum?id=ECvgmYVyeUz
AUTHORS: Yifei Wang, Qi Zhang, Yisen Wang, Jiansheng Yang, Zhouchen Lin
HIGHLIGHT: In this paper, we propose a new guarantee on the downstream performance without resort to the conditional independence assumption that is widely adopted in previous work but hardly holds in practice.

313, TITLE: Deep ReLU Networks Preserve Expected Length
https://openreview.net/forum?id=ci7LBzDn2Q
AUTHORS: Boris Hanin, Ryan S Jeong, David Rolnick
HIGHLIGHT: This article proves that, both on average and with high probability, randomly initialized ReLU networks with width larger than depth do not distort lengths and volumes.

314, TITLE: On Evaluation Metrics for Graph Generative Models
https://openreview.net/forum?id=EnwCZixjSh
AUTHORS: Rylee Thompson, Boris Knyazev, Elahe Ghalebi, Jungtaek Kim, Graham W. Taylor
HIGHLIGHT: In this work, we mitigate these issues by finding \emph{a scalar, domain-agnostic, and scalable metric} for evaluating and ranking GGMs.

315, TITLE: Bregman Gradient Policy Optimization
https://openreview.net/forum?id=ZU-zFnTum1N
AUTHORS: Feihu Huang, Shangqian Gao, Heng Huang
HIGHLIGHT: In this paper, we design a novel Bregman gradient policy optimization framework for reinforcement learning based on Bregman divergences and momentum techniques.

316, TITLE: Recursive Disentanglement Network
https://openreview.net/forum?id=CSfcOznpDY
AUTHORS: Yixuan Chen, Yubin Shi, Dongsheng Li, Yujiang Wang, Mingzhi Dong, Yingying Zhao, Robert Dick, Qin Lv, Fan Yang, Li Shang
HIGHLIGHT: This paper has described a solution to the compositional disentangled representation learning problem.

317, TITLE: PAC Prediction Sets Under Covariate Shift
https://openreview.net/forum?id=DhP9L8vIyLc
AUTHORS: Sangdon Park, Edgar Dobriban, Insup Lee, Osbert Bastani
HIGHLIGHT: We propose a novel algorithm that constructs a prediction set with probably approximated correct (PAC) guarantee under covariate shift, while minimizing the expected prediction set size.

318, TITLE: SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations
https://openreview.net/forum?id=aBsCjcPu_tE
AUTHORS: Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, Stefano Ermon
HIGHLIGHT: To address these issues, we introduce a new image synthesis and editing method, Stochastic Differential Editing (SDEdit), based on a diffusion model generative prior, which synthesizes realistic images by iteratively denoising through a stochastic differential equation (SDE).

319, TITLE: PolyLoss: A Polynomial Expansion Perspective of Classification Loss Functions
https://openreview.net/forum?id=gSdSJoenupI
AUTHORS: Zhaoqi Leng, Mingxing Tan, Chenxi Liu, Ekin Dogus Cubuk, Jay Shi, Shuyang Cheng, Dragomir Anguelov

HIGHLIGHT: In the PolyLoss framework, we propose a simple and effective Poly-1 formulation which outperforms the cross-entropy loss and focal loss on various of tasks.

320, TITLE: Fairness Guarantees under Demographic Shift
https://openreview.net/forum?id=wbPObLm6ueA
AUTHORS: Stephen Giguere, Blossom Metevier, Yuriy Brun, Philip S. Thomas, Scott Niekum, Bruno Castro da Silva
HIGHLIGHT: We propose a strategy for designing classification algorithms that provide high-confidence fairness guarantees that remain valid if the distribution over observations changes after the trained model is deployed.

321, TITLE: No Parameters Left Behind: Sensitivity Guided Adaptive Learning Rate for Training Large Transformer Models
https://openreview.net/forum?id=cuvga_CiVND
AUTHORS: Chen Liang, Haoming Jiang, Simiao Zuo, Pengcheng He, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, Tuo Zhao
HIGHLIGHT: We propose a novel adaptive learning rate training strategy for large Transformer models that encourages all parameters to be trained sufficiently.

322, TITLE: Learning Curves for Gaussian Process Regression with Power-Law Priors and Targets
https://openreview.net/forum?id=KeI9E-gsoB
AUTHORS: Hui Jin, Pradeep Kr. Banerjee, Guido Montufar
HIGHLIGHT: We derive the power-law decay rate of the generalization error in Gaussian process regression depending on the eigenspectrum of the prior and the target.

323, TITLE: Bi-linear Value Networks for Multi-goal Reinforcement Learning
https://openreview.net/forum?id=LedObtLmCjS
AUTHORS: Ge Yang, Zhang-Wei Hong, Pulkit Agrawal
HIGHLIGHT: We propose a bilinear value function for multi-goal reinforcement learning and show superior sample efficiency and generalizability.

324, TITLE: Frequency-aware SGD for Efficient Embedding Learning with Provable Benefits
https://openreview.net/forum?id=ibqTBNfJmi
AUTHORS: Yan Li, Dhruv Choudhary, Xiaohan Wei, Baichuan Yuan, Bhargav Bhushanam, Tuo Zhao, Guanghui Lan
HIGHLIGHT: Specifically, we propose (Counter-based) Frequency-aware Stochastic Gradient Descent, which applies a frequency-dependent learning rate for each token, and exhibits provable speed-up compared to SGD when the token distribution is imbalanced.

325, TITLE: ComPhy: Compositional Physical Reasoning of Objects and Events from Videos
https://openreview.net/forum?id=PgNEYaIc81Q
AUTHORS: Zhenfang Chen, Kexin Yi, Yunzhu Li, Mingyu Ding, Antonio Torralba, Joshua B. Tenenbaum, Chuang Gan
HIGHLIGHT: We introduce a new dataset for Compositional Physical Reasoning

326, TITLE: Efficient Token Mixing for Transformers via Adaptive Fourier Neural Operators
https://openreview.net/forum?id=EXHG-A3jlM
AUTHORS: John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, Bryan Catanzaro
HIGHLIGHT: We propose Adaptive Fourier Neural Operators (AFNO) for scaling self-attention to high resolution images in vision transformers by establishing a link between operator learning and token mixing.

327, TITLE: Learning Temporally Latent Causal Processes from General Temporal Data
https://openreview.net/forum?id=RDlLMjLJXdq
AUTHORS: Weiran Yao, Yuewen Sun, Alex Ho, Changyin Sun, Kun Zhang
HIGHLIGHT: Propose two provable conditions and training framework with which temporally latent causal processes are identifiable from observed variables.

328, TITLE: CROP: Certifying Robust Policies for Reinforcement Learning through Functional Smoothing
https://openreview.net/forum?id=HOjLHrlZhmx
AUTHORS: Fan Wu, Linyi Li, Zijian Huang, Yevgeniy Vorobeychik, Ding Zhao, Bo Li
HIGHLIGHT: We propose two particular types of robustness certification criteria: robustness of per-state actions and lower bound of cumulative rewards.

329, TITLE: MonoDistill: Learning Spatial Features for Monocular 3D Object Detection

https://openreview.net/forum?id=C54V-xTWfi
AUTHORS: Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, Wanli Ouyang
HIGHLIGHT: We propose the MonoDistill, which introduces spatial cues to the monocular 3D detector based on the knowledge distillation mechanism.

330, TITLE: Acceleration of Federated Learning with Alleviated Forgetting in Local Training
https://openreview.net/forum?id=541PxiEKN3F
AUTHORS: Chencheng Xu, Zhiwei Hong, Minlie Huang, Tao Jiang
HIGHLIGHT: Here, we propose FedReg, an algorithm to accelerate FL with alleviated knowledge forgetting in the local training stage by regularizing locally trained parameters with the loss on generated pseudo data, which encode the knowledge of previous training data learned by the global model.

331, TITLE: Curriculum learning as a tool to uncover learning principles in the brain
https://openreview.net/forum?id=TpJMvo0_pu-
AUTHORS: Daniel R. Kepple, Rainer Engelken, Kanaka Rajan
HIGHLIGHT: We present a novel approach to use curricula to identify principles by which a system learns.

332, TITLE: Large Learning Rate Tames Homogeneity: Convergence and Balancing Effect
https://openreview.net/forum?id=3tbDrs77LJ5
AUTHORS: Yuqing Wang, Minshuo Chen, Tuo Zhao, Molei Tao
HIGHLIGHT: Large learning rate well beyond 2/L provably induces an implicit regularization effect of balancing in gradient descent for matrix factorization.

333, TITLE: Granger causal inference on DAGs identifies genomic loci regulating transcription
https://openreview.net/forum?id=nZOUYEN6Wvy
AUTHORS: Alexander P Wu, Rohit Singh, Bonnie Berger
HIGHLIGHT: We show how to extend Granger causality to DAG-structured dynamical systems using graph neural networks, applying it to infer noncoding regions involved in gene regulation.

334, TITLE: Contrastive Learning is Just Meta-Learning
https://openreview.net/forum?id=gICys3ITSmj
AUTHORS: Renkun Ni, Manli Shu, Hossein Souri, Micah Goldblum, Tom Goldstein
HIGHLIGHT: We discuss the close relationship between contrastive learning and meta-learning, and we propose a meta-learning framework for self-supervised learning (SSL) along with meta-specific methods to improve contrastive learning performance for SSL.

335, TITLE: An Agnostic Approach to Federated Learning with Class Imbalance
https://openreview.net/forum?id=Xo0lbDt975
AUTHORS: Zebang Shen, Juan Cervino, Hamed Hassani, Alejandro Ribeiro
HIGHLIGHT: In this paper we propose a novel agnostic constrained learning formulation to tackle the class imbalance problem in FL, without requiring further information beyond the standard FL objective.

336, TITLE: Equivariant and Stable Positional Encoding for More Powerful Graph Neural Networks
https://openreview.net/forum?id=e95i1IHcWj
AUTHORS: Haorui Wang, Haoteng Yin, Muhan Zhang, Pan Li
HIGHLIGHT: In this work, we revisit GNNs that allow using positional features of nodes given by positional encoding (PE) techniques such as Laplacian Eigenmap, Deepwalk, etc..

337, TITLE: Invariant Causal Representation Learning for Out-of-Distribution Generalization
https://openreview.net/forum?id=-e4EXDWXnSn
AUTHORS: Chaochao Lu, Yuhuai Wu, Jos? Miguel Hern?ndez-Lobato, Bernhard Sch?lkopf
HIGHLIGHT: We propose invariant Causal Representation Learning (iCaRL), an approach that enables out-of-distribution (OOD) generalization in the nonlinear setting (i.e., nonlinear representations and nonlinear classifiers).

338, TITLE: Sound and Complete Neural Network Repair with Minimality and Locality Guarantees
https://openreview.net/forum?id=xS8AMYiEav3
AUTHORS: Feisi Fu, Wenchao Li
HIGHLIGHT: We present a novel methodology for repairing neural networks that use ReLU activation functions.

339, TITLE: Provable Learning-based Algorithm For Sparse Recovery

https://openreview.net/forum?id=BwPaPxwgyQb
AUTHORS: Xinshi Chen, Haoran Sun, Le Song
HIGHLIGHT: In this work, we propose PLISA (Provable Learning-based Iterative Sparse recovery Algorithm) to learn algorithms automatically from data.

340, TITLE: Learning curves for continual learning in neural networks: Self-knowledge transfer and forgetting
https://openreview.net/forum?id=tFgdrQbbaa
AUTHORS: Ryo Karakida, Shotaro Akaho
HIGHLIGHT: We analyze the generalization performance of continual learning in the NTK regime and identify key properties of knowledge transfer and forgetting.

341, TITLE: Reinforcement Learning in Presence of Discrete Markovian Context Evolution
https://openreview.net/forum?id=CmsfC7u054S
AUTHORS: Hang Ren, Aivar Sootla, Taher Jafferjee, Junxiao Shen, Jun Wang, Haitham Bou Ammar
HIGHLIGHT: We consider a context-dependent Reinforcement Learning (RL) setting, which is characterized by: a) an unknown finite number of not directly observable contexts; b) abrupt (discontinuous) context changes; and c) Markovian context evolution.

342, TITLE: Conditional Image Generation by Conditioning Variational Auto-Encoders
https://openreview.net/forum?id=7MV6uLzOChW
AUTHORS: William Harvey, Saeid Naderiparizi, Frank Wood
HIGHLIGHT: We create fast-to-train conditional VAEs using amortized inference in pretrained unconditional VAEs, and demonstrate diverse samples on image completion tasks.

343, TITLE: Neural Energy Minimization for Molecular Conformation Optimization
https://openreview.net/forum?id=7QfLW-XZTl
AUTHORS: Jiaqi Guan, Wesley Wei Qian, qiang liu, Wei-Ying Ma, Jianzhu Ma, Jian Peng
HIGHLIGHT: We propose a neural energy minimization formulation that casts the prediction problem into an unrolled optimization process, where a neural network is parametrized to learn the gradient fields of a conformational energy landscape.

344, TITLE: CoMPS: Continual Meta Policy Search
https://openreview.net/forum?id=PVJ6j87gOHz
AUTHORS: Glen Berseth, Zhiwei Zhang, Grace Zhang, Chelsea Finn, Sergey Levine
HIGHLIGHT: Cotinual meta-reinforcement learning accelerates task learning, via repeated meta off-policy search.

345, TITLE: Optimizing Few-Step Diffusion Samplers by Gradient Descent
https://openreview.net/forum?id=VFBjuF8HEp
AUTHORS: Daniel Watson, William Chan, Jonathan Ho, Mohammad Norouzi
HIGHLIGHT: We propose a method to discover fast, high-fidelity samplers for diffusion probabilistic models.

346, TITLE: The Evolution of Uncertainty of Learning in Games
https://openreview.net/forum?id=Fza94Y8VS4a
AUTHORS: Yun Kuen Cheung, Georgios Piliouras, Yixin Tao
HIGHLIGHT: We show that the differential entropy of certain learning-in-game systems increases linearly with time, formalizing their increased unpredictability over time.

347, TITLE: An Experimental Design Perspective on Exploration in Reinforcement Learning
https://openreview.net/forum?id=0no8Motr-zO
AUTHORS: Viraj Mehta, Biswajit Paria, Jeff Schneider, Willie Neiswanger, Stefano Ermon
HIGHLIGHT: We draw a connection between Bayesian Optimal Experiment Design and RL to develop an acquisition function to guide data collection in model based RL leading to improved sample efficiency.

348, TITLE: Revisiting flow generative models for Out-of-distribution detection
https://openreview.net/forum?id=6y2KBh-0Fd9
AUTHORS: Dihong Jiang, Sun Sun, Yaoliang Yu
HIGHLIGHT: In this work, we aim to re-examine the potential of generative flow models in OOD detection.

349, TITLE: On feature learning in shallow and multi-layer neural networks with global convergence guarantees
https://openreview.net/forum?id=PQTW3iG4sC-
AUTHORS: Zhengdao Chen, Eric Vanden-Eijnden, Joan Bruna

HIGHLIGHT: Gradient flow can induce feature learning in shallow and multi-layer neural networks while admitting non-asymptotic guarantees of global convergence.

350, TITLE: Differentiable Scaffolding Tree for Molecule Optimization
https://openreview.net/forum?id=w_drCosT76
AUTHORS: Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W. Coley, Jimeng Sun
HIGHLIGHT: make the molecular optimization problem differentiable at the structure level

351, TITLE: Skill-based Meta-Reinforcement Learning
https://openreview.net/forum?id=jeLW-Fh9bV
AUTHORS: Taewook Nam, Shao-Hua Sun, Karl Pertsch, Sung Ju Hwang, Joseph J Lim
HIGHLIGHT: In this work, we devise a method that enables meta-learning on long-horizon, sparse-reward tasks, allowing us to solve unseen target tasks with orders of magnitude fewer environment interactions.

352, TITLE: Shuffle Private Stochastic Convex Optimization
https://openreview.net/forum?id=DrZXuTGg2A-
AUTHORS: Albert Cheu, Matthew Joseph, Jieming Mao, Binghui Peng
HIGHLIGHT: The first analysis of shuffle private stochastic convex optimization.

353, TITLE: Fast topological clustering with Wasserstein distance
https://openreview.net/forum?id=0kPL3xO4R5
AUTHORS: Tananun Songdechakraiwut, Bryan M Krause, Matthew I Banks, Kirill V Nourski, Barry D Van Veen
HIGHLIGHT: In this paper, we propose a novel and computationally practical topological clustering method that clusters complex networks with intricate topology using principled theory from persistent homology and optimal transport.

354, TITLE: Learning Audio-Visual Speech Representation by Masked Multimodal Cluster Prediction
https://openreview.net/forum?id=Z1Qlm11uOM
AUTHORS: Bowen Shi, Wei-Ning Hsu, Kushal Lakhotia, Abdelrahman Mohamed
HIGHLIGHT: A self-supervised learning framework for audio-visual speech data, which uses only 30h of labeled data to match the SOTA lip-reading model trained on 31k hours of data (34.6% vs 33.6% WER), and further outperforms the SOTA with 70x less data (30.6%).

355, TITLE: Distilling GANs with Style-Mixed Triplets for X2I Translation with Limited Data
https://openreview.net/forum?id=QjOQkpzKbNk
AUTHORS: Yaxing Wang, Joost van de weijer, Lu Yu, SHANGLING JUI
HIGHLIGHT: One transfer learning method generalizes varying kinds of conditional image synthesization tasks.

356, TITLE: DEGREE: Decomposition Based Explanation for Graph Neural Networks
https://openreview.net/forum?id=Ve0Wth3ptT_
AUTHORS: Qizhang Feng, Ninghao Liu, Fan Yang, Ruixiang Tang, Mengnan Du, Xia Hu
HIGHLIGHT: We propose a new decomposition based explanation for Graph Neural Networks.

357, TITLE: GDA-AM: ON THE EFFECTIVENESS OF SOLVING MIN-IMAX OPTIMIZATION VIA ANDERSON MIXING
https://openreview.net/forum?id=3YqeuCVwy1d
AUTHORS: Huan He, Shifan Zhao, Yuanzhe Xi, Joyce Ho, Yousef Saad
HIGHLIGHT: We propose a new minimax optimization framework,GDA-AM, that views the GDA dynamics as a fixed-point iteration and solves it using Anderson Mixing to converge to the local minimax.

358, TITLE: Approximation and Learning with Deep Convolutional Models: a Kernel Perspective
https://openreview.net/forum?id=lrocYB-0ST2
AUTHORS: Alberto Bietti
HIGHLIGHT: We study the inductive bias of multi-layer convolutional models through a kernel lens, showing generalization benefits of various architectural choices such as locality, depth, and pooling layers.

359, TITLE: Synchromesh: Reliable Code Generation from Pre-trained Language Models
https://openreview.net/forum?id=KmtVD97J43e
AUTHORS: Gabriel Poesia, Alex Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, Sumit Gulwani
HIGHLIGHT: A framework to generate programs from large pre-trained language models (e.g. GPT-3, Codex) while satisfying syntactic and semantic constraints.

360, TITLE: Do deep networks transfer invariances across classes?
https://openreview.net/forum?id=Fn7i_r5rR0q
AUTHORS: Allan Zhou, Fahim Tajwar, Alexander Robey, Tom Knowles, George J. Pappas, Hamed Hassani, Chelsea Finn
HIGHLIGHT: Study how well classifiers learn invariances in the imbalanced setting, and methods for improvement.

361, TITLE: The Spectral Bias of Polynomial Neural Networks
https://openreview.net/forum?id=P7FLfMLTSEX
AUTHORS: Moulik Choraria, Leello Tadesse Dadi, Grigorios Chrysos, Julien Mairal, Volkan Cevher
HIGHLIGHT: We study the spectral bias of polynomial networks and compare it with the spectral bias of standard neural nets using kernel approximations

362, TITLE: Graph-less Neural Networks: Teaching Old MLPs New Tricks Via Distillation
https://openreview.net/forum?id=4p6_5HBWPCw
AUTHORS: Shichang Zhang, Yozen Liu, Yizhou Sun, Neil Shah
HIGHLIGHT: Distill knowledge from GNNs to MLPs to accelerate model inference and facilitate deployment for large-scale applications

363, TITLE: Triangle and Four Cycle Counting with Predictions in Graph Streams
https://openreview.net/forum?id=8in_5gN9I0
AUTHORS: Justin Y Chen, Talya Eden, Piotr Indyk, Honghao Lin, Shyam Narayanan, Ronitt Rubinfeld, Sandeep Silwal, Tal Wagner, David Woodruff, Michael Zhang
HIGHLIGHT: We propose algorithms for counting triangles and four cycles in graph streams with the aid of ML predictors.

364, TITLE: Object Pursuit: Building a Space of Objects via Discriminative Weight Generation
https://openreview.net/forum?id=lbauk6wK2-y
AUTHORS: Chuanyu Pan, Yanchao Yang, Kaichun Mo, Yueqi Duan, Leonidas Guibas
HIGHLIGHT: We propose a novel framework named object pursuit that can continuously learn object-centric representations using training data collected from interactions with individual objects.

365, TITLE: Learning Neural Contextual Bandits through Perturbed Rewards
https://openreview.net/forum?id=7inCJ3MhXt3
AUTHORS: Yiling Jia, Weitong ZHANG, Dongruo Zhou, Quanquan Gu, Hongning Wang
HIGHLIGHT: We propose to perturb the rewards when updating the neural network to eliminate the need of explicit exploration and the corresponding computational overhead.

366, TITLE: Defending Against Image Corruptions Through Adversarial Augmentations
https://openreview.net/forum?id=jJOjjiZHy3h
AUTHORS: Dan Andrei Calian, Florian Stimberg, Olivia Wiles, Sylvestre-Alvise Rebuffi, Andr?s Gy?rgy, Timothy A Mann, Sven Gowal
HIGHLIGHT: Our theoretically-supported method finds adversarial examples by optimizing over the weights of pre-trained autoencoders, and yields classifiers with improved robustness to image corruptions.

367, TITLE: Sample and Computation Redistribution for Efficient Face Detection
https://openreview.net/forum?id=RhB1AdoFfGE
AUTHORS: Jia Guo, Jiankang Deng, Alexandros Lattas, Stefanos Zafeiriou
HIGHLIGHT: We search for optimised computation distribution and training sample distribution for the task of face detection.

368, TITLE: Space-Time Graph Neural Networks
https://openreview.net/forum?id=XJiajt89Omg
AUTHORS: Samar Hadou, Charilaos I Kanatsoulis, Alejandro Ribeiro
HIGHLIGHT: We introduce space-time graph neural network (ST-GNN) tailored to jointly process the underlying space-time topology of time-varying network data, and we show its stability to perturbations.

369, TITLE: A First-Occupancy Representation for Reinforcement Learning
https://openreview.net/forum?id=JBAZe2yN6Ub
AUTHORS: Ted Moskovitz, Spencer R Wilson, Maneesh Sahani
HIGHLIGHT: We introduce the first-occupancy representation, a modification of the successor representation which enables agents to perform rapid policy evaluation and planning for a class of ethologically important non-Markovian reward functions.

370, TITLE: Towards Better Understanding and Better Generalization of Low-shot Classification in Histology Images with Contrastive Learning
https://openreview.net/forum?id=kQ2SOflIOVC
AUTHORS: Jiawei Yang, Hanbo Chen, Jiangpeng Yan, Xiaoyu Chen, Jianhua Yao
HIGHLIGHT: Here, we pioneer the study of low-shot learning for histology images by setting up three cross-domain tasks that reflect real clinics problems.

371, TITLE: QDrop: Randomly Dropping Quantization for Extremely Low-bit Post-Training Quantization
https://openreview.net/forum?id=ySQH0oDyp7
AUTHORS: Xiuying Wei, Ruihao Gong, Yuhang Li, Xianglong Liu, Fengwei Yu
HIGHLIGHT: In this study, we pioneeringly confirm that properly incorporating activation quantization into the PTQ reconstruction benefits the final accuracy.

372, TITLE: Policy Smoothing for Provably Robust Reinforcement Learning
https://openreview.net/forum?id=mwdfai8NBrJ
AUTHORS: Aounon Kumar, Alexander Levine, Soheil Feizi
HIGHLIGHT: A provable adversarial robustness technique for reinforcement learning.

373, TITLE: Learning Prototype-oriented Set Representations for Meta-Learning
https://openreview.net/forum?id=WH6u2SvlLp4
AUTHORS: Dan dan Guo, Long Tian, Minghe Zhang, Mingyuan Zhou, Hongyuan Zha
HIGHLIGHT: A plug-and-play framework for set-structured input tasks

374, TITLE: Attention-based Interpretability with Concept Transformers
https://openreview.net/forum?id=kAa9eDS0RdO
AUTHORS: Mattia Rigotti, Christoph Miksovic, Ioana Giurgiu, Thomas Gschwind, Paolo Scotton
HIGHLIGHT: We the Concept Transformer an architecture that generalization attention from low-level input features to high-level concepts as a mechanism to ensure the interpretability of attention scores within a given application domain.

375, TITLE: Wisdom of Committees: An Overlooked Approach To Faster and More Accurate Models
https://openreview.net/forum?id=MvO2t0vbs4-
AUTHORS: Xiaofang Wang, Dan Kondratyuk, Eric Christiansen, Kris M. Kitani, Yair Movshovitz-Attias, Elad Eban
HIGHLIGHT: A simple ensemble or cascade of off-the-shelf pre-trained models can match or exceed the accuracy of SOTA models while being drastically more efficient.

376, TITLE: Neural Network Approximation based on Hausdorff distance of Zonotopes
https://openreview.net/forum?id=oiZJwC_fyS
AUTHORS: Panagiotis Misiakos, Georgios Smyrnis, George Retsinas, Petros Maragos
HIGHLIGHT: Based on this result, we propose geometrical neural network compression methods that employ the K-means algorithm.

377, TITLE: Provably convergent quasistatic dynamics for mean-field two-player zero-sum games
https://openreview.net/forum?id=MP904TiHqJ-
AUTHORS: Chao Ma, Lexing Ying
HIGHLIGHT: We propose a quasistatic Wasserstein flow for finding mixed Nash equilibriums, and prove its convergence.

378, TITLE: DriPP: Driven Point Processes to Model Stimuli Induced Patterns in M/EEG Signals
https://openreview.net/forum?id=d_2lcDh0Y9c
AUTHORS: C?dric Allain, Alexandre Gramfort, Thomas Moreau
HIGHLIGHT: Model for patterns' activation using temporal point processes to reveal stimulus-induced effects in brain electrophysiology.

379, TITLE: Selective Ensembles for Consistent Predictions
https://openreview.net/forum?id=HfUyCRBeQc
AUTHORS: Emily Black, Klas Leino, Matt Fredrikson
HIGHLIGHT: Deep models give inconsistent predictions and explanations over small changes (e.g. random initialization). We can mitigate this by using selective ensemble models, which abstain from prediction if their constituent models do not agree sufficiently.

380, TITLE: Tuformer: Data-Driven Design of Expressive Transformer by Tucker Tensor Representation
https://openreview.net/forum?id=V0A5g83gdQ_
AUTHORS: Xiaoyu Liu, Jiahao Su, Furong Huang
HIGHLIGHT: We propose Tuformer, a data-driven design of theoretically guaranteed expressive Transformer with trainable heads, inspired by Tucker tensor representation.

381, TITLE: A Fine-Tuning Approach to Belief State Modeling
https://openreview.net/forum?id=ckZY7DGa7FQ
AUTHORS: Samuel Sokota, Hengyuan Hu, David J Wu, J Zico Kolter, Jakob Nicolaus Foerster, Noam Brown
HIGHLIGHT: A method for improving the accuracy of belief state models and for approximating public belief states at scale.

382, TITLE: Rethinking Network Design and Local Geometry in Point Cloud: A Simple Residual MLP Framework
https://openreview.net/forum?id=3Pbra-_u76D
AUTHORS: Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, Yun Fu
HIGHLIGHT: In this paper, we present a new design for point cloud analysis , dubbed as PointMLP, which delivers new state-of-the-art results on multiple benchmarks and exhibits gratifying inference speed.

383, TITLE: Neural Methods for Logical Reasoning over Knowledge Graphs
https://openreview.net/forum?id=tgcAoUVHRIB
AUTHORS: Alfonso Amayuelas, Shuai Zhang, Xi Susie Rao, Ce Zhang
HIGHLIGHT: Neural Network models for answering multi-hop queris on Knowledge Graphs
We introduce a set of models that use Neural Networks to create one-point vector embeddings to answer the queries.

384, TITLE: Explainable GNN-Based Models over Knowledge Graphs
https://openreview.net/forum?id=CrCvGNHAIrz
AUTHORS: David Jaime Tena Cucala, Bernardo Cuenca Grau, Egor V. Kostylev, Boris Motik
HIGHLIGHT: We propose a new family of graph neural network-based transformations of graph data that can be trained effectively and where all predictions can be explained symbolically as logical inferences in Datalog.

385, TITLE: Hierarchical Few-Shot Imitation with Skill Transition Models
https://openreview.net/forum?id=xKZ4K0lTj_
AUTHORS: Kourosh Hakhamaneshi, Ruihan Zhao, Albert Zhan, Pieter Abbeel, Michael Laskin
HIGHLIGHT: We introduce a new algorithm (FIST) which extracts skills from offline data and adapts them in few-shot to solve unseen complex long-horizon tasks by utilizing an inverse skill dynamics model and semi-parametric imitation.

386, TITLE: It Takes Two to Tango: Mixup for Deep Metric Learning
https://openreview.net/forum?id=ZKy2X3dgPA
AUTHORS: Shashanka Venkataramanan, Bill Psomas, Ewa Kijak, laurent amsaleg, Konstantinos Karantzalos, Yannis Avrithis
HIGHLIGHT: We systematically study different mixup strategies in the context of deep metric learning, and study how to, and the effect of, mixing inputs, features, and output embeddings.

387, TITLE: Hindsight: Posterior-guided training of retrievers for improved open-ended generation
https://openreview.net/forum?id=Vr_BTpw3wz
AUTHORS: Ashwin Paranjape, Omar Khattab, Christopher Potts, Matei Zaharia, Christopher D Manning
HIGHLIGHT: We use a posterior-guide retriever to train a retrieval-augmented generation that performs well on open-ended one-to-many generation tasks.

388, TITLE: Learning Transferable Reward for Query Object Localization with Policy Adaptation
https://openreview.net/forum?id=92tYQiil17
AUTHORS: Tingfeng Li, Shaobo Han, Martin Renqiang Min, Dimitris N. Metaxas
HIGHLIGHT: We propose a reinforcement learning based approach to the problem of query object localization, where an agent is trained to localize objects of interest specified by a small exemplary set.

389, TITLE: One After Another: Learning Incremental Skills for a Changing World
https://openreview.net/forum?id=dg79moSRqIo
AUTHORS: Nur Muhammad Mahi Shafiullah, Lerrel Pinto
HIGHLIGHT: We discover skills incrementally, without supervision, and it lets us learn skills in both static environments and environments with changing dynamics.

390, TITLE: Heteroscedastic Temporal Variational Autoencoder For Irregularly Sampled Time Series
https://openreview.net/forum?id=Az7opqbQE-3
AUTHORS: Satya Narayan Shukla, Benjamin Marlin
HIGHLIGHT: We present a new deep learning architecture for probabilistic interpolation of irregularly sampled time series.

391, TITLE: CodeTrek: Flexible Modeling of Code using an Extensible Relational Representation
https://openreview.net/forum?id=WQc075jmBmf
AUTHORS: Pardis Pashakhanloo, Aaditya Naik, Yuepeng Wang, Hanjun Dai, Petros Maniatis, Mayur Naik
HIGHLIGHT: We present a relational database representation and corresponding neural module for source code and show its potential on program understanding tasks

392, TITLE: How to deal with missing data in supervised deep learning?
https://openreview.net/forum?id=J7b4BCtDm4
AUTHORS: Niels Bruun Ipsen, Pierre-Alexandre Mattei, Jes Frellsen
HIGHLIGHT: Marginalize over missing values in supervised learning using deep latent variable models.

393, TITLE: Local Feature Swapping for Generalization in Reinforcement Learning
https://openreview.net/forum?id=Sq0-tgDyHe4
AUTHORS: David Bertoin, Emmanuel Rachelson
HIGHLIGHT: We propose a simple yet effective layer increasing the generalization abilities of reinforcement learning agents

394, TITLE: Practical Conditional Neural Process Via Tractable Dependent Predictions
https://openreview.net/forum?id=3pugbNqOh5m
AUTHORS: Stratis Markou, James Requeima, Wessel Bruinsma, Anna Vaughan, Richard E Turner
HIGHLIGHT: In this work, we present a new class of Neural Process models that produce correlated predictions and support exact maximum likelihood training that is simple and scalable.

395, TITLE: Value Function Spaces: Skill-Centric State Abstractions for Long-Horizon Reasoning
https://openreview.net/forum?id=vgqS1vkkCbE
AUTHORS: Dhruv Shah, Alexander T Toshev, Sergey Levine, brian ichter
HIGHLIGHT: We introduce value function spaces, a learned representation of state through the values of low-level skills, which capture affordances and ignores distractors to enable long-horizon reasoning and zero-shot generalization.

396, TITLE: Transfer RL across Observation Feature Spaces via Model-Based Regularization
https://openreview.net/forum?id=7KdAoOsI81C
AUTHORS: Yanchao Sun, Ruijie Zheng, Xiyao Wang, Andrew E Cohen, Furong Huang
HIGHLIGHT: We propose a model-based transfer learning algorithm that transfers knowledge across tasks with drastically different observation spaces, without any prior knowledge of the inter-task mapping.

397, TITLE: Symbolic Learning to Optimize: Towards Interpretability and Scalability
https://openreview.net/forum?id=ef0nInZHKIC
AUTHORS: Wenqing Zheng, Tianlong Chen, Ting-Kuei Hu, Zhangyang Wang
HIGHLIGHT: Learning to distill learned optimization rule into symbolic math equations that bears better interpretability and scales better.

398, TITLE: The Uncanny Similarity of Recurrence and Depth
https://openreview.net/forum?id=3wNcr5nq56
AUTHORS: Avi Schwarzschild, Arjun Gupta, Amin Ghiasi, Micah Goldblum, Tom Goldstein
HIGHLIGHT: We show quantitatively and qualitatively that recurrent models have the same behaviors as feed-forward networks despite reusing parameters at each recurrence.

399, TITLE: Diurnal or Nocturnal? Federated Learning of Multi-branch Networks from Periodically Shifting Distributions
https://openreview.net/forum?id=E4EE_ohFGz
AUTHORS: Chen Zhu, Zheng Xu, Mingqing Chen, Jakub Konecn?, Andrew Hard, Tom Goldstein
HIGHLIGHT: We study a better modeling assumption for the periodical distribution shift in FL, and propose algorithms that learn better from the shifting distribution.

400, TITLE: Relating transformers to models and neural representations of the hippocampal formation
https://openreview.net/forum?id=B8DVo9B1YE0
AUTHORS: James C. R. Whittington, Joseph Warren, Tim E.J. Behrens

HIGHLIGHT: Transformers learn brain representatations and they are algorithmically related to models of the hippocampal formation.

401, TITLE: ADAVI: Automatic Dual Amortized Variational Inference Applied To Pyramidal Bayesian Models
https://openreview.net/forum?id=CgIEctmcXx1
AUTHORS: Louis Rouillard, Demian Wassermann
HIGHLIGHT: We automatically derive a variational family dual to a plate-enriched Hierarchical Bayesian Network and perform amortized inference.

402, TITLE: Trans-Encoder: Unsupervised sentence-pair modelling through self- and mutual-distillations
https://openreview.net/forum?id=AmUhwTOHgm
AUTHORS: Fangyu Liu, Yunlong Jiao, Jordan Massiah, Emine Yilmaz, Serhii Havrylov
HIGHLIGHT: Bootstrapping an unsupervised sentence encoder by self-distilling knowledge between its bi-encoder and cross-encoder forms, enhancing each other iteratively.

403, TITLE: Group equivariant neural posterior estimation
https://openreview.net/forum?id=u6s8dSporO8
AUTHORS: Maximilian Dax, Stephen R Green, Jonathan Gair, Michael Deistler, Bernhard Sch?lkopf, Jakob H. Macke
HIGHLIGHT: We describe a method to incorporate group equivariances into neural posterior estimation.

404, TITLE: Finding an Unsupervised Image Segmenter in each of your Deep Generative Models
https://openreview.net/forum?id=Ug-bgjgSlKV
AUTHORS: Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, Andrea Vedaldi
HIGHLIGHT: We propose an entirely unsupervised method for foreground-background image segmentation based on automatically finding a direction in the latent space of a deep generative model.

405, TITLE: Generalization Through the Lens of Leave-One-Out Error
https://openreview.net/forum?id=7grkzyj89A_
AUTHORS: Gregor Bachmann, Thomas Hofmann, Aurelien Lucchi
HIGHLIGHT: In this work, we instead revisit the concept of leave-one-out error to measure the generalization ability of deep and wide models.

406, TITLE: No One Representation to Rule Them All: Overlapping Features of Training Methods
https://openreview.net/forum?id=BK-4qbGgIE3
AUTHORS: Raphael Gontijo-Lopes, Yann Dauphin, Ekin Dogus Cubuk
HIGHLIGHT: We study the effect of training methodology on prediction diversity and show that diverging training setups produce diverse features, uncorrelated errors, and more efficient ensembles.

407, TITLE: Charformer: Fast Character Transformers via Gradient-based Subword Tokenization
https://openreview.net/forum?id=JtBRnrlOEFN
AUTHORS: Yi Tay, Vinh Q. Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, Donald Metzler
HIGHLIGHT: Fast Token-Free Models

408, TITLE: On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks
https://openreview.net/forum?id=aPOpXlnV1T
AUTHORS: Maximilian Seitzer, Arash Tavakoli, Dimitrije Antic, Georg Martius
HIGHLIGHT: We analyse problems with the training objective of probabilistic neural networks and propose a fix in the form of a new loss function.

409, TITLE: Variational Neural Cellular Automata
https://openreview.net/forum?id=7fFO4cMBx_9
AUTHORS: Rasmus Berg Palm, Miguel Gonz?lez Duque, Shyam Sudhakaran, Sebastian Risi
HIGHLIGHT: We propose and evaluate the Variational Neural Cellular Automata, a self-organising generative model based on neural cellular automata

410, TITLE: PSA-GAN: Progressive Self Attention GANs for Synthetic Time Series
https://openreview.net/forum?id=Ix_mh42xq5w
AUTHORS: Paul Jeha, Michael Bohlke-Schneider, Pedro Mercado, Shubham Kapoor, Rajbir Singh Nirwan, Valentin Flunkert, Jan Gasthaus, Tim Januschowski

HIGHLIGHT: In this paper we present PSA-GAN, a generative adversarial network (GAN) that generates long time series samples of high quality using progressive growing of GANs and self-attention.

411, TITLE: Prospect Pruning: Finding Trainable Weights at Initialization using Meta-Gradients
https://openreview.net/forum?id=AIgn9uwfcD1
AUTHORS: Milad Alizadeh, Shyam A. Tailor, Luisa M Zintgraf, Joost van Amersfoort, Sebastian Farquhar, Nicholas Donald Lane, Yarin Gal
HIGHLIGHT: We use meta-gradients to prune neural networks at initialization based on "trainability" of weights instead of their impact on the loss at a single step.

412, TITLE: Learning 3D Representations of Molecular Chirality with Invariance to Bond Rotations
https://openreview.net/forum?id=hm2tNDdgaFK
AUTHORS: Keir Adams, Lagnajit Pattanaik, Connor W. Coley
HIGHLIGHT: We propose a method of processing the 3D torsion angles of a molecular conformer to learn tetrahedral chirality while integrating a novel invariance to rotations about internal molecular bonds directly into the model architecture.

413, TITLE: SketchODE: Learning neural sketch representation in continuous time
https://openreview.net/forum?id=c-4HSDAWua5
AUTHORS: Ayan Das, Yongxin Yang, Timothy Hospedales, Tao Xiang, Yi-Zhe Song
HIGHLIGHT: Modelling continuous time chirographic structures like handwriting, diagrams, sketches etc with Neural Ordinary Differential Equations.

414, TITLE: Efficient Active Search for Combinatorial Optimization Problems
https://openreview.net/forum?id=nO5caZwFwYu
AUTHORS: Andr? Hottung, Yeong-Dae Kwon, Kevin Tierney
HIGHLIGHT: We propose active search approaches for combinatorial optimization problems that search for solutions by adjusting a subset of (model) parameters to a single instance at test time.

415, TITLE: Generalized rectifier wavelet covariance models for texture synthesis
https://openreview.net/forum?id=ziRLU3Y2PN_
AUTHORS: Antoine Brochard, Sixin Zhang
HIGHLIGHT: This paper presents a model for texture synthesis, built on a wavelet-based representation of images.

416, TITLE: FlexConv: Continuous Kernel Convolutions With Differentiable Kernel Sizes
https://openreview.net/forum?id=3jooF27-0Wy
AUTHORS: David W. Romero, Robert-Jan Bruintjes, Jakub Mikolaj Tomczak, Erik J Bekkers, Mark Hoogendoorn, Jan van Gemert
HIGHLIGHT: We provide a high bandwidth, alias-free convolutional kernel parameterization with learnable kernel size and constant parameter cost.

417, TITLE: Collapse by Conditioning: Training Class-conditional GANs with Limited Data
https://openreview.net/forum?id=7TZeCsNOUB_
AUTHORS: Mohamad Shahbazi, Martin Danelljan, Danda Pani Paudel, Luc Van Gool
HIGHLIGHT: Motivated by this observation, we propose a training strategy for conditional GANs (cGANs) that effectively prevents the observed mode-collapse by leveraging unconditional learning.

418, TITLE: Bootstrapping Semantic Segmentation with Regional Contrast
https://openreview.net/forum?id=6u6N8WWwYSM
AUTHORS: Shikun Liu, Shuaifeng Zhi, Edward Johns, Andrew Davison
HIGHLIGHT: We present a pixel-level contrastive learning framework to achieve a high-quality semantic segmeantation model trained with very few human annotations.

419, TITLE: Capacity of Group-invariant Linear Readouts from Equivariant Representations: How Many Objects can be Linearly Classified Under All Possible Views?
https://openreview.net/forum?id=_4GFbtOuWq-
AUTHORS: Matthew Farrell, Blake Bordelon, Shubhendu Trivedi, Cengiz Pehlevan
HIGHLIGHT: We compute the fraction of linearly separable dichotomies of representations formed by group actions and apply these results to intermediate representations of convolutional neural networks.

420, TITLE: Using Graph Representation Learning with Schema Encoders to Measure the Severity of Depressive Symptoms

https://openreview.net/forum?id=OtEDS2NWhqa
AUTHORS: Simin Hong, Anthony Cohn, David Crossland Hogg
HIGHLIGHT: We encode word embeddings by using graph representation learning method, which schematizes context-level depressive features for depression state prediction.

421, TITLE: A Program to Build E(N)-Equivariant Steerable CNNs
https://openreview.net/forum?id=WE4qe9xlnQw
AUTHORS: Gabriele Cesa, Leon Lang, Maurice Weiler
HIGHLIGHT: We derive a general method to build G-steerable kernel spaces for equivariant steerable CNNs

422, TITLE: Why Propagate Alone? Parallel Use of Labels and Features on Graphs
https://openreview.net/forum?id=VTNjxbFRKly
AUTHORS: Yangkun Wang, Jiarui Jin, Weinan Zhang, Yang Yongyi, Jiuhai Chen, Quan Gan, Yong Yu, Zheng Zhang, Zengfeng Huang, David Wipf
HIGHLIGHT: We analyze a powerful label trick from a theoretical perspective and reduce it to an interpretable form that inspires broader application scenarios.

423, TITLE: A global convergence theory for deep ReLU implicit networks via over-parameterization
https://openreview.net/forum?id=R332S76RjxS
AUTHORS: Tianxiang Gao, Hailiang Liu, Jia Liu, Hridesh Rajan, Hongyang Gao
HIGHLIGHT: For a ReLU activated implicit neural network with infinitely many layers, we prove that randomly initialized gradient descent with a fixed step-size converges to a global minimum at a linear rate if the width m is the squared sample size n.

424, TITLE: TRAIL: Near-Optimal Imitation Learning with Suboptimal Data
https://openreview.net/forum?id=6q_2b6u0BnJ
AUTHORS: Mengjiao Yang, Sergey Levine, Ofir Nachum
HIGHLIGHT: A provably beneficial way to learn action representations for imitation learning from suboptimal auxiliary data.

425, TITLE: Mind the Gap: Domain Gap Control for Single Shot Domain Adaptation for Generative Adversarial Networks
https://openreview.net/forum?id=vqGi8Kp0wM
AUTHORS: Peihao Zhu, Rameen Abdal, John Femiani, Peter Wonka
HIGHLIGHT: We propose several regularizers to control the domain transfer for single shot domain adaptation in the context of generative adversarial networks.

426, TITLE: Gradient Step Denoiser for convergent Plug-and-Play
https://openreview.net/forum?id=fPhKeld3Okz
AUTHORS: Samuel Hurault, Arthur Leclaire, Nicolas Papadakis
HIGHLIGHT: We propose a performant Plug-and-Play image restoration algorithm that theoretically converges with an exact gradient step deep denoiser.

427, TITLE: Deep Point Cloud Reconstruction
https://openreview.net/forum?id=mKDtUtxIGJ
AUTHORS: Jaesung Choe, ByeongIn Joung, Francois Rameau, Jaesik Park, In So Kweon
HIGHLIGHT: We propose deep learning-based point cloud reconstruction algorithm

428, TITLE: Critical Points in Quantum Generative Models
https://openreview.net/forum?id=2f1z55GVQN
AUTHORS: Eric Ricardo Anschuetz
HIGHLIGHT: We show using techniques from random matrix theory that, unlike typical neural networks, quantum generative models often have poor quality local minima.

429, TITLE: Clean Images are Hard to Reblur: Exploiting the Ill-Posed Inverse Task for Dynamic Scene Deblurring
https://openreview.net/forum?id=kezNJydWvE
AUTHORS: Seungjun Nah, Sanghyun Son, Jaerin Lee, Kyoung Mu Lee
HIGHLIGHT: Reblurring, the inverse task of deblurring, is used for supervised/self-supervised learning of deblurring and improves the image sharpness.

430, TITLE: Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift
https://openreview.net/forum?id=cGDAkQo1C0p
AUTHORS: Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, Jaegul Choo

HIGHLIGHT: We propose a simple yet effective normalization method, reversible instance normalization (RevIN), which solves the time-series forecasting task against the distribution shift problem.

431, TITLE: Reducing Excessive Margin to Achieve a Better Accuracy vs. Robustness Trade-off
https://openreview.net/forum?id=Azh9QBQ4tR7
AUTHORS: Rahul Rade, Seyed-Mohsen Moosavi-Dezfooli
HIGHLIGHT: In this paper, we closely examine the changes induced in the decision boundary of a deep network during adversarial training.

432, TITLE: Maximum Entropy RL (Provably) Solves Some Robust RL Problems
https://openreview.net/forum?id=PtSAD3caaA2
AUTHORS: Benjamin Eysenbach, Sergey Levine
HIGHLIGHT: Maximum Entropy RL (provably) solves some robust RL problems.

433, TITLE: Learning Continuous Environment Fields via Implicit Functions
https://openreview.net/forum?id=3ILxkQ7yElm
AUTHORS: Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Ming-Hsuan Yang, Jan Kautz
HIGHLIGHT: We propose a novel scene representation that can dynamically change behaviors of agents inside the scene.

434, TITLE: Learning Distributionally Robust Models at Scale via Composite Optimization
https://openreview.net/forum?id=To-R742x7se
AUTHORS: Farzin Haddadpour, Mohammad Mahdi Kamani, Mehrdad Mahdavi, amin karbasi
HIGHLIGHT: In this paper, we show how different variants of DRO are simply instances of a finite-sum composite optimization for which we provide scalable methods.

435, TITLE: EXACT: Scalable Graph Neural Networks Training via Extreme Activation Compression
https://openreview.net/forum?id=vkaMaq95_rX
AUTHORS: Zirui Liu, Kaixiong Zhou, Fan Yang, Li Li, Rui Chen, Xia Hu
HIGHLIGHT: Based on the implementation, we propose a memory-efficient framework called ``EXACT", which for the first time demonstrate the potential and evaluate the feasibility of training GNNs with compressed activations.

436, TITLE: VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning
https://openreview.net/forum?id=xm6YD62D1Ub
AUTHORS: Adrien Bardes, Jean Ponce, Yann LeCun
HIGHLIGHT: Variance regularization prevents collapse in self-supervised representation learning

437, TITLE: Amortized Implicit Differentiation for Stochastic Bilevel Optimization
https://openreview.net/forum?id=3PN4iyXBeF
AUTHORS: Michael Arbel, Julien Mairal
HIGHLIGHT: We provide a unified framework for analyzing bilevel optimization algorithm based on approximate implicit differentiation with a warm-start strategy

438, TITLE: Model Agnostic Interpretability for Multiple Instance Learning
https://openreview.net/forum?id=KSSfF5lMIAg
AUTHORS: Joseph Early, Christine Evers, SArvapali Ramchurn
HIGHLIGHT: We propose and compare several methods for model-agnostic interpretability for multiple instance learning.

439, TITLE: Non-Parallel Text Style Transfer with Self-Parallel Supervision
https://openreview.net/forum?id=-TSe5o7STVR
AUTHORS: Ruibo Liu, Chongyang Gao, Chenyan Jia, Guangxuan Xu, Soroush Vosoughi
HIGHLIGHT: We propose a new text style transfer model for non-parallel corpus with supervision from intrinsic parallelism.

440, TITLE: Towards Understanding the Robustness Against Evasion Attack on Categorical Data
https://openreview.net/forum?id=BmJV7kyAmg
AUTHORS: Hongyan Bao, Yufei Han, Yujun Zhou, Yun Shen, Xiangliang Zhang
HIGHLIGHT: This paper explores the characterization and certification of the robustness against evasion attack with categorical input.

441, TITLE: Anti-Concentrated Confidence Bonuses For Scalable Exploration

https://openreview.net/forum?id=RXQ-FPbQYVn
AUTHORS: Jordan T. Ash, Cyril Zhang, Surbhi Goel, Akshay Krishnamurthy, Sham M. Kakade
HIGHLIGHT: We introduce anti-concentrated confidence bounds for efficiently approximating the elliptical bonus, using an ensemble of regressors trained to predict random noise from policy network-derived features.

442, TITLE: Constrained Physical-Statistics Models for Dynamical System Identification and Prediction
https://openreview.net/forum?id=gbe1zHyA73
AUTHORS: J?r?mie DONA, Marie D?chelle, patrick gallinari, Marina Levy
HIGHLIGHT: We propose to incorporate constraints in the learning of hybrid physical and data driven dynamical models.

443, TITLE: Understanding Dimensional Collapse in Contrastive Self-supervised Learning
https://openreview.net/forum?id=YevsQ05DEN7
AUTHORS: Li Jing, Pascal Vincent, Yann LeCun, Yuandong Tian
HIGHLIGHT: We observe and theoretically explain the dimensional collapse in contrastive self-supervised learning. We also propose a novel SSL method that does not rely on a projector.

444, TITLE: THOMAS: Trajectory Heatmap Output with learned Multi-Agent Sampling
https://openreview.net/forum?id=QDdJhACYrlX
AUTHORS: Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanciulescu, Fabien Moutarde
HIGHLIGHT: We propose a solution for multi-agent coherent multimodal trajectory prediction by learning a recombination of each agent predicted modalities.

445, TITLE: Group-based Interleaved Pipeline Parallelism for Large-scale DNN Training
https://openreview.net/forum?id=cw-EmNq5zfD
AUTHORS: PengCheng Yang, Xiaoming Zhang, Wenpeng Zhang, Ming Yang, Hong Wei
HIGHLIGHT: In this work, we propose a novel pipeline named WPipe, which achieves better memory efficiency and fresher weight updates.

446, TITLE: HyAR: Addressing Discrete-Continuous Action Reinforcement Learning via Hybrid Action Representation
https://openreview.net/forum?id=64trBbOhdGU
AUTHORS: Boyan Li, Hongyao Tang, YAN ZHENG, Jianye HAO, Pengyi Li, Zhen Wang, Zhaopeng Meng, LI Wang
HIGHLIGHT: In this paper, we propose Hybrid Action Representation (HyAR) to learn a compact and decodable latent representation space for the original hybrid action space.

447, TITLE: Constrained Graph Mechanics Networks
https://openreview.net/forum?id=SHbhHHfePhP
AUTHORS: Wenbing Huang, Jiaqi Han, Yu Rong, Tingyang Xu, Fuchun Sun, Junzhou Huang
HIGHLIGHT: In this paper, we propose Graph Mechanics Network (GMN) which is combinatorially efficient, equivariant and constraint-aware.

448, TITLE: Zero Pixel Directional Boundary by Vector Transform
https://openreview.net/forum?id=nxcABL7jbQh
AUTHORS: Edoardo Mello Rella, Ajad Chhatkuli, Yun Liu, Ender Konukoglu, Luc Van Gool
HIGHLIGHT: In this paper, we re-interpret boundaries as 1-D surfaces and formulate a one-to-one vector transform function that allows for training of boundary prediction completely avoiding the class imbalance issue.

449, TITLE: Bayesian Modeling and Uncertainty Quantification for Learning to Optimize: What, Why, and How
https://openreview.net/forum?id=EVVadRFRgL7
AUTHORS: Yuning You, Yue Cao, Tianlong Chen, Zhangyang Wang, Yang Shen
HIGHLIGHT: Our study represents the first effort to recognize and quantify the uncertainty of the optimization algorithm.

450, TITLE: Online Continual Learning on Class Incremental Blurry Task Configuration with Anytime Inference
https://openreview.net/forum?id=nrGGfMbY_qK
AUTHORS: Hyunseo Koh, Dahyun Kim, Jung-Woo Ha, Jonghyun Choi
HIGHLIGHT: a novel continual learning set-up that is online and task-free, has new class distributions and focuses on any-time inference

451, TITLE: Towards Training Billion Parameter Graph Neural Networks for Atomic Simulations
https://openreview.net/forum?id=0jP2n0YFmKG
AUTHORS: Anuroop Sriram, Abhishek Das, Brandon M Wood, C. Lawrence Zitnick

HIGHLIGHT: We scale GNNs used for modeling atomic simulations by an order of magnitude and obtain large performance improvements on the Open Catalyst 2020 dataset.

452, TITLE: TAda! Temporally-Adaptive Convolutions for Video Understanding
https://openreview.net/forum?id=izj68lUcBpt
AUTHORS: Ziyuan Huang, Shiwei Zhang, Liang Pan, Zhiwu Qing, Mingqian Tang, Ziwei Liu, Marcelo H Ang Jr
HIGHLIGHT: A stand-alone temporal modelling module or a plug-in enhancement of the 1D/2D/3D convolutions used in video models for better and more efficient temporal modelling.

453, TITLE: You Mostly Walk Alone: Analyzing Feature Attribution in Trajectory Prediction
https://openreview.net/forum?id=POxF-LEqnF
AUTHORS: Osama Makansi, Julius Von K?gelgen, Francesco Locatello, Peter Vincent Gehler, Dominik Janzing, Thomas Brox, Bernhard Sch?lkopf
HIGHLIGHT: We propose a Shapley value-based method for attributing trajectory prediction performance to different input features and show on common benchmark datasets that existing models do not use interaction information, contrary to their claims.

454, TITLE: PEARL: Data Synthesis via Private Embeddings and Adversarial Reconstruction Learning
https://openreview.net/forum?id=M6M8BEmd6dq
AUTHORS: Seng Pei Liew, Tsubasa Takahashi, Michihiko Ueno
HIGHLIGHT: We propose a new framework of synthesizing data using deep generative models in a differentially private manner.

455, TITLE: Boosted Curriculum Reinforcement Learning
https://openreview.net/forum?id=anbBFlX1tJ1
AUTHORS: Pascal Klink, Carlo D'Eramo, Jan Peters, Joni Pajarinen
HIGHLIGHT: A novel approach for curriculum RL that increases the representativeness of the functional space as new, increasingly complex, tasks from the curriculum are presented to the agent.

456, TITLE: T-WaveNet: A Tree-Structured Wavelet Neural Network for Time Series Signal Analysis
https://openreview.net/forum?id=U4uFaLyg7PV
AUTHORS: Minhao LIU, Ailing Zeng, Qiuxia LAI, Ruiyuan Gao, Min Li, Jing Qin, Qiang Xu
HIGHLIGHT: In this work, we propose a novel tree-structured wavelet neural network for time series signal analysis, namely T-WaveNet, by taking advantage of an inherent property of various types of signals, known as the dominant frequency range.

457, TITLE: Cross-Domain Lossy Compression as Optimal Transport with an Entropy Bottleneck
https://openreview.net/forum?id=BRFWxcZfAdC
AUTHORS: Huan Liu, George Zhang, Jun Chen, Ashish J Khisti
HIGHLIGHT: We consider the novel task of cross-distribution lossy compression and characterize it as an optimal transport problem under an entropy constraint, then provide experimental results to demonstrate the principles suggested by our theory.

458, TITLE: Counterfactual Plans under Distributional Ambiguity
https://openreview.net/forum?id=noaG7SrPVK0
AUTHORS: Ngoc Bui, Duy Nguyen, Viet Anh Nguyen
HIGHLIGHT: This study the counterfactual plans under model uncertainty, in which the distribution of the model parameters is partially prescribed using only the first- and second-moment information.

459, TITLE: Energy-Based Learning for Cooperative Games, with Applications to Valuation Problems in Machine Learning
https://openreview.net/forum?id=xLfAgCroImw
AUTHORS: Yatao Bian, Yu Rong, Tingyang Xu, Jiaxiang Wu, Andreas Krause, Junzhou Huang
HIGHLIGHT: An energy-based treatment for cooperative games provides a decoupling perspective for Shapley value and others.

460, TITLE: Iterated Reasoning with Mutual Information in Cooperative and Byzantine Decentralized Teaming
https://openreview.net/forum?id=giBFoa-uS12
AUTHORS: Sachin G Konan, Esmaeil Seraj, Matthew Gombolay
HIGHLIGHT: We propose a MARL framework with iterated and hierarchical rationalizability through mutual information for decision-making in fully-decentralized cooperative and Byzantine scenarios.

461, TITLE: Constraining Linear-chain CRFs to Regular Languages
https://openreview.net/forum?id=jbrgwbv8nD

AUTHORS: Sean Papay, Roman Klinger, Sebastian Pado
HIGHLIGHT: CRFs can be efficiently constrained to arbitrary regular languages, enforcing nonlocal constraints on their outputs.

462, TITLE: Language-driven Semantic Segmentation
https://openreview.net/forum?id=RriDjddCLN
AUTHORS: Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, Rene Ranftl
HIGHLIGHT: We present a language-driven approach that enables synthesis of zero-shot semantic segmentation models from arbitrary label sets at test time.

463, TITLE: Towards General Function Approximation in Zero-Sum Markov Games
https://openreview.net/forum?id=sA4qIu3zv6v
AUTHORS: Baihe Huang, Jason D. Lee, Zhaoran Wang, Zhuoran Yang
HIGHLIGHT: In the decoupled setting where the agent controls a single player and plays against an arbitrary opponent, we propose a new model-free algorithm.

464, TITLE: Parallel Training of GRU Networks with a Multi-Grid Solver for Long Sequences
https://openreview.net/forum?id=N1WI0vJLER
AUTHORS: Euhyun Moon, Eric C Cyr
HIGHLIGHT: This paper presents a novel parallel-in-time training scheme for GRU networks based on a MGRIT solver.

465, TITLE: A Statistical Framework for Efficient Out of Distribution Detection in Deep Neural Networks
https://openreview.net/forum?id=Oy9WeuZD51
AUTHORS: Matan Haroush, Tzviel Frostig, Ruth Heller, Daniel Soudry
HIGHLIGHT: We frame Out Of Distribution (OOD) detection in DNNs as a statistical hypothesis testing problem.

466, TITLE: Visual Representation Learning Does Not Generalize Strongly Within the Same Domain
https://openreview.net/forum?id=9RUHPlladgh
AUTHORS: Lukas Schott, Julius Von K?gelgen, Frederik Tr?uble, Peter Vincent Gehler, Chris Russell, Matthias Bethge, Bernhard Sch?lkopf, Francesco Locatello, Wieland Brendel
HIGHLIGHT: We study and benchmark the inductive biases for generalization in visual representation learning on systematic out-of-distribution settings.

467, TITLE: Robust and Scalable SDE Learning: A Functional Perspective
https://openreview.net/forum?id=xZ6H7wydGl
AUTHORS: Scott Alexander Cameron, Tyron Luke Cameron, Arnu Pretorius, Stephen J. Roberts
HIGHLIGHT: We provide an algorithm for estimating the probability of observations of a stochastic process which is significantly faster and more stable than those based on standard integration schemes

468, TITLE: Equivariant Self-Supervised Learning: Encouraging Equivariance in Representations
https://openreview.net/forum?id=gKLAAfiytI
AUTHORS: Rumen Dangovski, Li Jing, Charlotte Loh, Seungwook Han, Akash Srivastava, Brian Cheung, Pulkit Agrawal, Marin Soljacic
HIGHLIGHT: Imposing invariance to certain transformations (e.g. random resized cropping) and sensitivity to other transformations (e.g. four-fold rotations) learns better features.

469, TITLE: L0-Sparse Canonical Correlation Analysis
https://openreview.net/forum?id=KntaNRo6R48
AUTHORS: Ofir Lindenbaum, Moshe Salhov, Amir Averbuch, Yuval Kluger
HIGHLIGHT: We propose a new ℓ_0-CCA method for learning correlated representations based on sparse subsets of variables from two observed modalities.

470, TITLE: Neural Processes with Stochastic Attention: Paying more attention to the context dataset
https://openreview.net/forum?id=JPkQwEdYn8
AUTHORS: Mingyu Kim, Kyeong Ryeol Go, Se-Young Yun
HIGHLIGHT: This paper extends the attentive neural process (ANP), replacing the deterministic weights in the cross-attention module of ANP with latent weights.

471, TITLE: Understanding and Improving Graph Injection Attack by Promoting Unnoticeability
https://openreview.net/forum?id=wkMG8cdvh7-

AUTHORS: Yongqiang Chen, Han Yang, Yonggang Zhang, MA KAILI, Tongliang Liu, Bo Han, James Cheng
HIGHLIGHT: We find GIA can be provably more harmful than GMA while at the price of bringing more damage to the homophily distribution, which makes it can easily defendable, hence we propose a novel adversarial objective to mitigate the issue.

472, TITLE: Continual Normalization: Rethinking Batch Normalization for Online Continual Learning
https://openreview.net/forum?id=vwLLQ-HwqhZ
AUTHORS: Quang Pham, Chenghao Liu, Steven HOI
HIGHLIGHT: A negative effect of BN in online continual learning and a simple strategy to alleviate it.

473, TITLE: Chunked Autoregressive GAN for Conditional Waveform Synthesis
https://openreview.net/forum?id=v3aeIsY_vVX
AUTHORS: Max Morrison, Rithesh Kumar, Kundan Kumar, Prem Seetharaman, Aaron Courville, Yoshua Bengio
HIGHLIGHT: We improve the state-of-the-art of conditional waveform synthesis by combining the strengths of GANs and autoregression

474, TITLE: Topological Graph Neural Networks
https://openreview.net/forum?id=oxxUMeFwEHd
AUTHORS: Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, Karsten Borgwardt
HIGHLIGHT: We describe a new layer for graph neural networks that incorporates multi-scale (ranging from local to global) topological information.

475, TITLE: Communication-Efficient Actor-Critic Methods for Homogeneous Markov Games
https://openreview.net/forum?id=xy_2w3J3kH
AUTHORS: Dingyang Chen, Yile Li, Qi Zhang
HIGHLIGHT: In this paper, we formally characterize a subclass of cooperative Markov games where agents exhibit a certain form of homogeneity such that policy sharing provably incurs no suboptimality.

476, TITLE: Toward Faithful Case-based Reasoning through Learning Prototypes in a Nearest Neighbor-friendly Space.
https://openreview.net/forum?id=R79ZGjHhv6p
AUTHORS: Seyed Omid Davoudi, Majid Komeili
HIGHLIGHT: Offering better prototype explanations using a nearest-neighbor friendly embedding space.

477, TITLE: Augmented Sliced Wasserstein Distances
https://openreview.net/forum?id=iMqTLyfwnOO
AUTHORS: Xiongjie Chen, Yongxin Yang, Yunpeng Li
HIGHLIGHT: In this work, we propose a new family of distance metrics, called augmented sliced Wasserstein distances (ASWDs), constructed by first mapping samples to higher-dimensional hypersurfaces parameterized by neural networks.

478, TITLE: cosFormer: Rethinking Softmax In Attention
https://openreview.net/forum?id=Bl8CQrx2Up4
AUTHORS: Qin Zhen, Weixuan Sun, Hui Deng, Dongxu Li, Yunshen Wei, Baohong Lv, Junjie Yan, Lingpeng Kong, Yiran Zhong
HIGHLIGHT: A new linear transformer.

479, TITLE: β-Intact-VAE: Identifying and Estimating Causal Effects under Limited Overlap
https://openreview.net/forum?id=q7n2RngwOM
AUTHORS: Pengzhou Abel Wu, Kenji Fukumizu
HIGHLIGHT: See all these naturally in one: limited overlap, prognostic score, identifiable VAE, balanced representation Learning, CATE error bounds.

480, TITLE: MaGNET: Uniform Sampling from Deep Generative Network Manifolds Without Retraining
https://openreview.net/forum?id=r5qumLiYwf9
AUTHORS: Ahmed Imtiaz Humayun, Randall Balestriero, Richard Baraniuk
HIGHLIGHT: We propose a differential-geometry-based technique to provably sample uniformly from the data manifold of a trained Deep Generative Network without the need for retraining.

481, TITLE: Multi-Agent MDP Homomorphic Networks
https://openreview.net/forum?id=H7HDG--DJF0
AUTHORS: Elise van der Pol, Herke van Hoof, Frans A Oliehoek, Max Welling
HIGHLIGHT: We introduce globally equivariant multi-agent policy networks with distributed execution.

482, TITLE: Bag of Instances Aggregation Boosts Self-supervised Distillation
https://openreview.net/forum?id=N0uJGWDw21d
AUTHORS: Haohang Xu, Jiemin Fang, XIAOPENG ZHANG, Lingxi Xie, Xinggang Wang, Wenrui Dai, Hongkai Xiong, Qi Tian
HIGHLIGHT: This paper proposes a new self-supervised distillation method which aggregates related instances bagged by the teacher and shows stronger performance than previous relation-agnostic methods.

483, TITLE: Quadtree Attention for Vision Transformers
https://openreview.net/forum?id=fR-EnKWL_Zb
AUTHORS: Shitao Tang, Jiahui Zhang, Siyu Zhu, Ping Tan
HIGHLIGHT: Reduce the computation of transformer by coarse-to-fine attention computation for feature matching/stereo/classification/detection tasks.

484, TITLE: Offline Reinforcement Learning for Large Scale Language Action Spaces
https://openreview.net/forum?id=qaxhBG1UUaS
AUTHORS: Youngsoo Jang, Jongmin Lee, Kee-Eung Kim
HIGHLIGHT: In this paper, we introduce GPT-Critic, an offline RL method for task-oriented dialogue.

485, TITLE: Adversarial Retriever-Ranker for Dense Text Retrieval
https://openreview.net/forum?id=MR7XubKUFB
AUTHORS: Hang Zhang, Yeyun Gong, Yelong Shen, Jiancheng Lv, Nan Duan, Weizhu Chen
HIGHLIGHT: To address these challenges, we present Adversarial Retriever-Ranker (AR2), which consists of a dual-encoder retriever plus a cross-encoder ranker.
We will make our code, models, and data publicly available.

486, TITLE: OBJECT DYNAMICS DISTILLATION FOR SCENE DECOMPOSITION AND REPRESENTATION
https://openreview.net/forum?id=oJGDYQFKL3i
AUTHORS: Qu Tang, Xiangyu Zhu, Zhen Lei, Zhaoxiang Zhang
HIGHLIGHT: In this paper, we work on object dynamics and propose Object Dynamics Distillation Network (ODDN), a framework that distillates explicit object dynamics (e.g., velocity) from sequential static representations.

487, TITLE: Igeood: An Information Geometry Approach to Out-of-Distribution Detection
https://openreview.net/forum?id=mfwdY3U_9ea
AUTHORS: Eduardo Dadalto Camara Gomes, Florence Alberge, Pierre Duhamel, Pablo Piantanida
HIGHLIGHT: We propose a flexible and effective out-of-distribution detection method by building on the Fisher-Rao distance between probability distributions.

488, TITLE: CoST: Contrastive Learning of Disentangled Seasonal-Trend Representations for Time Series Forecasting
https://openreview.net/forum?id=PilZY3omXV2
AUTHORS: Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, Steven Hoi
HIGHLIGHT: Following this principle, we propose a new time series representation learning framework for time series forecasting named CoST, which applies contrastive learning methods to learn disentangled seasonal-trend representations.

489, TITLE: Information Bottleneck: Exact Analysis of (Quantized) Neural Networks
https://openreview.net/forum?id=kF9DZQQrU0w
AUTHORS: Stephan Sloth Lorenzen, Christian Igel, Mads Nielsen
HIGHLIGHT: We investigate the information bottleneck in quantized neural networks, allowing us to compute the exact mutual information and provide an analysis free from estimation artifacts.

490, TITLE: Relational Surrogate Loss Learning
https://openreview.net/forum?id=dZPgfwaTaXv
AUTHORS: Tao Huang, Zekang Li, Hua Lu, Yong Shan, Shusheng Yang, Yang Feng, Fei Wang, Shan You, Chang Xu
HIGHLIGHT: In this paper, we show that directly maintaining the relation of models between surrogate losses and metrics suffices, and propose a rank correlation-based optimization method to maximize this relation and learn surrogate losses.

491, TITLE: Goal-Directed Planning via Hindsight Experience Replay
https://openreview.net/forum?id=6NePxZwfae
AUTHORS: Lorenzo Moro, Amarildo Likmeta, Enrico Prati, Marcello Restelli
HIGHLIGHT: This paper presents an extension of AlphaZero to tackle sparse reward goal-based tasks

492, TITLE: Learning Generalizable Representations for Reinforcement Learning via Adaptive Meta-learner of Behavioral Similarities
https://openreview.net/forum?id=zBOI9LFpESK
AUTHORS: Jianda Chen, Sinno Pan
HIGHLIGHT: In this paper, we propose a novel meta-learner-based framework for representation learning regarding behavioral similarities in reinforcement learning.

493, TITLE: Filling the G_ap_s: Multivariate Time Series Imputation by Graph Neural Networks
https://openreview.net/forum?id=kOu3-S3wJ7
AUTHORS: Andrea Cini, Ivan Marisca, Cesare Alippi
HIGHLIGHT: We propose a graph neural network architecture for multivariate time series imputation and achieve state-of-the-art results on several benchmarks.

494, TITLE: On the Existence of Universal Lottery Tickets
https://openreview.net/forum?id=SYB4WrJql1n
AUTHORS: Rebekka Burkholz, Nilanjana Laha, Rajarshi Mukherjee, Alkis Gotovos
HIGHLIGHT: We formalize a notion of strong universal lottery tickets and prove their existence as subnetworks of randomly initialized neural networks.

495, TITLE: Understanding the Variance Collapse of SVGD in High Dimensions
https://openreview.net/forum?id=Qycd9j5Qp9J
AUTHORS: Jimmy Ba, Murat A Erdogdu, Marzyeh Ghassemi, Shengyang Sun, Taiji Suzuki, Denny Wu, Tianzong Zhang
HIGHLIGHT: Qualitative and quantitative analysis of the variance collapse phenomenon of SVGD in high dimensions.

496, TITLE: Learning to Annotate Part Segmentation with Gradient Matching
https://openreview.net/forum?id=zNR43c03lRy
AUTHORS: Yu Yang, Xiaotian Cheng, Hakan Bilen, Xiangyang Ji
HIGHLIGHT: We propose a gradient-matching-based method to learn annotator which is able to label generated images with part segmentation by decoding the generator features into segmentation masks.

497, TITLE: Temporal Efficient Training of Spiking Neural Network via Gradient Re-weighting
https://openreview.net/forum?id=_XNtisL32jv
AUTHORS: Shikuang Deng, Yuhang Li, Shanghang Zhang, Shi Gu
HIGHLIGHT: This paper provides a novel temporal efficient training method for SNN, which significantly improves performance by modifying the optimization target.

498, TITLE: Graphon based Clustering and Testing of Networks: Algorithms and Theory
https://openreview.net/forum?id=sTNHCrIKDQc
AUTHORS: Mahalakshmi Sabanayagam, Leena Chennuru Vankadara, Debarghya Ghoshdastidar
HIGHLIGHT: In this work, we propose methods for clustering multiple graphs, without vertex correspondence, that are inspired by the recent literature on estimating graphons---symmetric functions corresponding to infinite vertex limit of graphs.

499, TITLE: Gradient Matching for Domain Generalization
https://openreview.net/forum?id=vDwBW49HmO
AUTHORS: Yuge Shi, Jeffrey Seely, Philip Torr, Siddharth N, Awni Hannun, Nicolas Usunier, Gabriel Synnaeve
HIGHLIGHT: We propose to learn features that are invariant across domains by maximizing the gradient inner product between domains.

500, TITLE: Discovering Latent Concepts Learned in BERT
https://openreview.net/forum?id=POTMtpYI1xH
AUTHORS: Fahim Dalvi, Abdul Rafae Khan, Firoj Alam, Nadir Durrani, Jia Xu, Hassan Sajjad
HIGHLIGHT: We approach interpretability from a model?s perspective by discovering and analyzing latent concepts learned in pre-trained models in an unsupervised fashion.
We also release a novel BERT concept dataset (BCD) consisting of 198 concept labels and 1M annotated instances.

501, TITLE: Generalized Demographic Parity for Group Fairness
https://openreview.net/forum?id=YigKlMJwjye
AUTHORS: Zhimeng Jiang, Xiaotian Han, Chao Fan, Fan Yang, Ali Mostafavi, Xia Hu

HIGHLIGHT: This work aims to generalize demographic parity to continuous sensitive attributes while preserving tractable computation.

502, TITLE: Representing Mixtures of Word Embeddings with Mixtures of Topic Embeddings
https://openreview.net/forum?id=IYMuTbGzjFU
AUTHORS: dongsheng wang, Dan dan Guo, He Zhao, Huangjie Zheng, Korawat Tanwisuth, Bo Chen, Mingyuan Zhou
HIGHLIGHT: A novel method to learn in the word embedding space a set of globally shared topic embedding vectors that are manifested differently in each document

503, TITLE: AdaAug: Learning Class- and Instance-adaptive Data Augmentation Policies
https://openreview.net/forum?id=rWXfFogxRJN
AUTHORS: Tsz-Him Cheung, Dit-Yan Yeung
HIGHLIGHT: We propose a novel Automated Data Augmentation method called \texttt{AdaAug} to efficiently learn adaptive augmentation policies in a class-dependent and potentially instance-dependent manner.

504, TITLE: Robust Unlearnable Examples: Protecting Data Privacy Against Adversarial Learning
https://openreview.net/forum?id=baUQQPwQiAg
AUTHORS: Shaopeng Fu, Fengxiang He, Yang Liu, Li Shen, Dacheng Tao
HIGHLIGHT: This paper proposes an robust error-minimizing noise that can protect data from being learned under adversarial training.

505, TITLE: Deep Attentive Variational Inference
https://openreview.net/forum?id=T4-65DNlDij
AUTHORS: Ifigeneia Apostolopoulou, Ian Char, Elan Rosenfeld, Artur Dubrawski
HIGHLIGHT: Specifically, we propose deep attentive variational autoencoder and test it on a variety of established datasets.

506, TITLE: Supervision Exists Everywhere: A Data Efficient Contrastive Language-Image Pre-training Paradigm
https://openreview.net/forum?id=zq1iJkNk3uN
AUTHORS: Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, Junjie Yan
HIGHLIGHT: This work proposes a novel training paradigm, Data efficient CLIP (DeCLIP), to alleviate this limitation.

507, TITLE: πBO: Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization
https://openreview.net/forum?id=MMAeCXIa89
AUTHORS: Carl Hvarfner, Danny Stoll, Artur Souza, Luigi Nardi, Marius Lindauer, Frank Hutter
HIGHLIGHT: We extend the Bayesian Optimization framework by allowing for arbitrary user priors over promising regions of the search space, to guide the search towards said regions.

508, TITLE: Topologically Regularized Data Embeddings
https://openreview.net/forum?id=P1QUVhOtEFP
AUTHORS: Robin Vandaele, Bo Kang, Jefrey Lijffijt, Tijl De Bie, Yvan Saeys
HIGHLIGHT: A method for incorporating expert prior topological knowledge into data embeddings.

509, TITLE: How Much Can CLIP Benefit Vision-and-Language Tasks?
https://openreview.net/forum?id=zf_Ll3HZWgy
AUTHORS: Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, Kurt Keutzer
HIGHLIGHT: To further study the advantage brought by CLIP, we propose to use CLIP as the visual encoder in various V&L models in two typical scenarios: 1) plugging CLIP into task-specific fine-tuning; 2) combining CLIP with V&L pre-training and transferring to downstream tasks.

510, TITLE: Fast Model Editing at Scale
https://openreview.net/forum?id=0DcZxeWfOPt
AUTHORS: Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, Christopher D Manning
HIGHLIGHT: A computationally efficient approach for learning to edit the behavior of very large pre-trained language models (10 billion+ parameters)

511, TITLE: Leveraging unlabeled data to predict out-of-distribution performance
https://openreview.net/forum?id=o_HsiMPYh_x
AUTHORS: Saurabh Garg, Sivaraman Balakrishnan, Zachary Chase Lipton, Behnam Neyshabur, Hanie Sedghi

HIGHLIGHT: In this work, we investigate methods for predicting the target domain accuracy using only labeled source data and unlabeled target data.

512, TITLE: Optimization and Adaptive Generalization of Three layer Neural Networks
https://openreview.net/forum?id=dPyRNUlttBv
AUTHORS: Khashayar Gatmiry, Stefanie Jegelka, Jonathan Kelner
HIGHLIGHT: Algorithmically obtaining noise-robust and adaptive generalization bounds for a three layer network model by going beyond the linear approximation of the network

513, TITLE: COPA: Certifying Robust Policies for Offline Reinforcement Learning against Poisoning Attacks
https://openreview.net/forum?id=psh0oeMSBiF
AUTHORS: Fan Wu, Linyi Li, Huan Zhang, Bhavya Kailkhura, Krishnaram Kenthapadi, Ding Zhao, Bo Li
HIGHLIGHT: We propose the first framework for certifiying robustness of offline reinforcement learning against poisoning attacks.

514, TITLE: HTLM: Hyper-Text Pre-Training and Prompting of Language Models
https://openreview.net/forum?id=P-pPW1nxf1r
AUTHORS: Armen Aghajanyan, Dmytro Okhonko, Mike Lewis, Mandar Joshi, Hu Xu, Gargi Ghosh, Luke Zettlemoyer
HIGHLIGHT: We unlock new state-of-the-art ways of priming and automatically generating prompts by pre-training on simplified HTML.

515, TITLE: Privacy Implications of Shuffling
https://openreview.net/forum?id=5i2f-aR6B8H
AUTHORS: Casey Meehan, Amrita Roy Chowdhury, Kamalika Chaudhuri, Somesh Jha
HIGHLIGHT: a novel formalization of the privacy offered by shuffling

516, TITLE: Image BERT Pre-training with Online Tokenizer
https://openreview.net/forum?id=ydopy-e6Dg
AUTHORS: Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, Tao Kong
HIGHLIGHT: We present a self-supervised framework iBOT that can perform masked image modeling with an online tokenizer, achieving the state-of-the-art results in downstream tasks.

517, TITLE: GNN is a Counter? Revisiting GNN for Question Answering
https://openreview.net/forum?id=hzmQ4wOnSb
AUTHORS: Kuan Wang, Yuyu Zhang, Diyi Yang, Le Song, Tao Qin
HIGHLIGHT: Counting is essential for reasoning and our simplistic graph neural counter is efficient and effective for QA tasks.

518, TITLE: FALCON: Fast Visual Concept Learning by Integrating Images, Linguistic descriptions, and Conceptual Relations
https://openreview.net/forum?id=htWIlvDcY8
AUTHORS: Lingjie Mei, Jiayuan Mao, Ziqi Wang, Chuang Gan, Joshua B. Tenenbaum
HIGHLIGHT: We present a meta-learning framework for learning new visual concepts quickly, from just one or a few examples, guided by multiple naturally occurring data streams: simultaneously looking at images, reading sentences that describe the objects in the scene, and interpreting supplemental sentences that relate the novel concept with other concepts.

519, TITLE: Toward Efficient Low-Precision Training: Data Format Optimization and Hysteresis Quantization
https://openreview.net/forum?id=3HJOA-1hb0e
AUTHORS: Sunwoo Lee, Jeongwoo Park, Dongsuk Jeon
HIGHLIGHT: We propose a systematic data format optimization method and hysteresis quantization scheme to enable efficient low-precision training.

520, TITLE: Discovering Nonlinear PDEs from Scarce Data with Physics-encoded Learning
https://openreview.net/forum?id=Vog_3GXsgmb
AUTHORS: Chengping Rao, Pu Ren, Yang Liu, Hao Sun
HIGHLIGHT: This work seeks to solve the data-driven governing equation discovery problem with a novel physics-encoded learning framework.

521, TITLE: Stiffness-aware neural network for learning Hamiltonian systems
https://openreview.net/forum?id=uVXEKeqJbNa

AUTHORS: SENWEI Liang, Zhongzhan Huang, Hong Zhang
HIGHLIGHT: We propose stiffness-aware neural network (SANN), a new method for learning Hamiltonian dynamical systems from data.

522, TITLE: InfinityGAN: Towards Infinite-Pixel Image Synthesis
https://openreview.net/forum?id=ufGMqIM0a4b
AUTHORS: Chieh Hubert Lin, Hsin-Ying Lee, Yen-Chi Cheng, Sergey Tulyakov, Ming-Hsuan Yang
HIGHLIGHT: InfinityGAN learns to synthesize arbitrary-sized images with limited resources and enables multiple new applications.

523, TITLE: Neural Contextual Bandits with Deep Representation and Shallow Exploration
https://openreview.net/forum?id=xnYACQquaGV
AUTHORS: Pan Xu, Zheng Wen, Handong Zhao, Quanquan Gu
HIGHLIGHT: A new neural network based algorithm for contextual bandit problems with theoretical guarantees and empirical advantages.

524, TITLE: SimVLM: Simple Visual Language Model Pretraining with Weak Supervision
https://openreview.net/forum?id=GUrhfTuf_3
AUTHORS: Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, Yuan Cao
HIGHLIGHT: In this work, we relax these constraints and present a minimalist pretraining framework, named Simple Visual Language Model (SimVLM).

525, TITLE: Gradient Importance Learning for Incomplete Observations
https://openreview.net/forum?id=fXHl76nO2AZ
AUTHORS: Qitong Gao, Dong Wang, Joshua David Amason, Siyang Yuan, Chenyang Tao, Ricardo Henao, Majda Hadziahmetovic, Lawrence Carin, Miroslav Pajic
HIGHLIGHT: In this work, we introduce the gradient importance learning (GIL) method to train multilayer perceptrons (MLPs) and long short-term memories (LSTMs) to directly perform inference from inputs containing missing values without imputation.

526, TITLE: Safe Neurosymbolic Learning with Differentiable Symbolic Execution
https://openreview.net/forum?id=NYBmJN4MyZ
AUTHORS: Chenxi Yang, Swarat Chaudhuri
HIGHLIGHT: We present DSE, the first approach to worst-case-safe parameter learning for potentially non-differentiable neurosymbolic programs where we bridge symbolic execution and stochastic gradient estimator to learn the loss of safety properties.

527, TITLE: Learning Weakly-supervised Contrastive Representations
https://openreview.net/forum?id=MSwEFaztwkE
AUTHORS: Yao-Hung Hubert Tsai, Tianqin Li, Weixin Liu, Peiyuan Liao, Ruslan Salakhutdinov, Louis-Philippe Morency
HIGHLIGHT: We present a weakly-supervised contrastive learning framework that considers auxiliary information (additional sources of information from data).

528, TITLE: Phenomenology of Double Descent in Finite-Width Neural Networks
https://openreview.net/forum?id=lTqGXfn9Tv
AUTHORS: Sidak Pal Singh, Aurelien Lucchi, Thomas Hofmann, Bernhard Sch?lkopf
HIGHLIGHT: We provide a theoretical analysis of double descent that applies for finite-width neural networks and delivers insights into the properties of neural networks (and its Hessian spectrum) near the interpolation threshold.

529, TITLE: Zero-Shot Self-Supervised Learning for MRI Reconstruction
https://openreview.net/forum?id=085y6YPaYjP
AUTHORS: Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Mehmet Akcakaya
HIGHLIGHT: Zero-shot self-supervised learning to perform robust subject-specific MRI reconstruction

530, TITLE: GiraffeDet: A Heavy-Neck Paradigm for Object Detection
https://openreview.net/forum?id=cBu4ElJfneV
AUTHORS: yiqi jiang, Zhiyu Tan, Junyan Wang, Xiuyu Sun, Ming Lin, Hao Li
HIGHLIGHT: we propose a novel heavy-neck paradigm(GiraffeDet) for detection task, which allows detectors to process the high-level categorical information and low-level spatial information uniformly, making it more effective in detection tasks.

531, TITLE: Pseudo Numerical Methods for Diffusion Models on Manifolds

https://openreview.net/forum?id=PlKWVd2yBkY
AUTHORS: Luping Liu, Yi Ren, Zhijie Lin, Zhou Zhao
HIGHLIGHT: We propose PNDMs, a new kind of numerical method, to accelerate diffusion models on manifolds.

532, TITLE: Doubly Adaptive Scaled Algorithm for Machine Learning Using Second-Order Information
https://openreview.net/forum?id=HCelXXcSEuH
AUTHORS: Majid Jahani, Sergey Rusakov, Zheng Shi, Peter Richt?rik, Michael W. Mahoney, Martin Takac
HIGHLIGHT: Second-Order Method for Large Scale Machine Learning Tasks

533, TITLE: Shallow and Deep Networks are Near-Optimal Approximators of Korobov Functions
https://openreview.net/forum?id=AV8FPoMTTa
AUTHORS: Moise Blanchard, Mohammed Amine Bennouna
HIGHLIGHT: We analyze the number of neurons and training parameters that a neural network needs to approximate multivariate functions of bounded second mixed derivatives

534, TITLE: Contextualized Scene Imagination for Generative Commonsense Reasoning
https://openreview.net/forum?id=Oh1r2wApbPv
AUTHORS: PeiFeng Wang, Jonathan Zamora, Junfeng Liu, Filip Ilievski, Muhao Chen, Xiang Ren
HIGHLIGHT: This work aims at tackling generative commonsense reasoning by allowing machines to imagine a reasonable scene before generating text.

535, TITLE: Efficient and Differentiable Conformal Prediction with General Function Classes
https://openreview.net/forum?id=Ht85_jyihxp
AUTHORS: Yu Bai, Song Mei, Huan Wang, Yingbo Zhou, Caiming Xiong
HIGHLIGHT: We generalize conformal prediction to learning multiple parameters within a general function class, to obtain an improved efficiency subject to valid coverage.

536, TITLE: Enhancing Cross-lingual Transfer by Manifold Mixup
https://openreview.net/forum?id=OjPmfr9GkVv
AUTHORS: Huiyun Yang, Huadong Chen, Hao Zhou, Lei Li
HIGHLIGHT: We propose the cross-lingual manifold mixup method to improve the cross-lingual transfer.

537, TITLE: Hybrid Memoised Wake-Sleep: Approximate Inference at the Discrete-Continuous Interface
https://openreview.net/forum?id=auOPcdAcoy
AUTHORS: Tuan Anh Le, Katherine M. Collins, Luke Hewitt, Kevin Ellis, Siddharth N, Samuel Gershman, Joshua B. Tenenbaum
HIGHLIGHT: Here, we propose Hybrid Memoised Wake-Sleep (HMWS), an algorithm for effective inference in such hybrid discrete-continuous models.

538, TITLE: Evading Adversarial Example Detection Defenses with Orthogonal Projected Gradient Descent
https://openreview.net/forum?id=af1eUDdUVz
AUTHORS: Oliver Bryniarski, Nabeel Hingun, Pedro Pachuca, Vincent Wang, Nicholas Carlini
HIGHLIGHT: We break four defenses that detect adversarial examples by introducing an improved attack technique that modifies the gradient before applying PGD.

539, TITLE: IFR-Explore: Learning Inter-object Functional Relationships in 3D Indoor Scenes
https://openreview.net/forum?id=OT3mLgR8Wg8
AUTHORS: QI LI, Kaichun Mo, Yanchao Yang, Hang Zhao, Leonidas Guibas
HIGHLIGHT: We formulate a novel problem of learning inter-object functional relationships in 3D indoor environments and propose a novel method that combines prior knowledge modeling and interactive policy learning to solve the task.
We create a new dataset based on the AI2Thor and PartNet datasets and perform extensive experiments that prove the effectiveness of our proposed method.

540, TITLE: Sound Adversarial Audio-Visual Navigation
https://openreview.net/forum?id=NkZq4OEYN-
AUTHORS: Yinfeng Yu, Wenbing Huang, Fuchun Sun, Changan Chen, Yikai Wang, Xiaohong Liu
HIGHLIGHT: This work aims to do an adversarial sound intervention for robust audio-visual navigation.

541, TITLE: Model-Based Offline Meta-Reinforcement Learning with Regularization
https://openreview.net/forum?id=EBn0uInJZWh

AUTHORS: Sen Lin, Jialin Wan, Tengyu Xu, Yingbin Liang, Junshan Zhang
HIGHLIGHT: This paper proposes a novel offline Meta-RL algorithm with regularization, which has provable performance improvement and outperforms the existing baselines empirically.

542, TITLE: On the Certified Robustness for Ensemble Models and Beyond
https://openreview.net/forum?id=tUa4REjGjTf
AUTHORS: Zhuolin Yang, Linyi Li, Xiaojun Xu, Bhavya Kailkhura, Tao Xie, Bo Li
HIGHLIGHT: Inspired by theoretical analysis, we propose Diversity Regularized Training to enhance the certified robustness of ensemble models and DRT significantly outperforms existing methods.

543, TITLE: Case-based Reasoning for Better Generalization in Text-Adventure Games
https://openreview.net/forum?id=ZDaSIkWT-AP
AUTHORS: Mattia Atzeni, Shehzaad Zuzar Dhuliawala, Keerthiram Murugesan, MRINMAYA SACHAN
HIGHLIGHT: In a departure from deep RL approaches, in this paper, we propose a general method inspired by case-based reasoning to train agents and generalize out of the training distribution.

544, TITLE: Network Augmentation for Tiny Deep Learning
https://openreview.net/forum?id=TYw3-OlrRm-
AUTHORS: Han Cai, Chuang Gan, Ji Lin, song han
HIGHLIGHT: a training method for tiny neural networks

545, TITLE: Meta-Imitation Learning by Watching Video Demonstrations
https://openreview.net/forum?id=KTPuIsx4pmo
AUTHORS: Jiayi Li, Tao Lu, Xiaoge Cao, Yinghao Cai, Shuo Wang
HIGHLIGHT: We present an approach of meta-imitation learning by watching video demonstrations from humans.

546, TITLE: C-Planning: An Automatic Curriculum for Learning Goal-Reaching Tasks
https://openreview.net/forum?id=K2JfSnLBD9
AUTHORS: Tianjun Zhang, Benjamin Eysenbach, Ruslan Salakhutdinov, Sergey Levine, Joseph E. Gonzalez
HIGHLIGHT: An algorithm for goal-conditioned RL that uses an automatic curriculum of waypoints during exploration, derived from variational inference.

547, TITLE: An Unconstrained Layer-Peeled Perspective on Neural Collapse
https://openreview.net/forum?id=WZ3yjh8coDg
AUTHORS: Wenlong Ji, Yiping Lu, Yiliang Zhang, Zhun Deng, Weijie J Su
HIGHLIGHT: We investigate how the gradient flow converges to a neural collapse solution in an unconstrained model.

548, TITLE: Hindsight Foresight Relabeling for Meta-Reinforcement Learning
https://openreview.net/forum?id=P7OVkHEoHOZ
AUTHORS: Michael Wan, Jian Peng, Tanmay Gangwani
HIGHLIGHT: We present HFR, a relabeling method that can be applied to meta-reinforcement learning to boost sample efficiency and performance.

549, TITLE: Optimizer Amalgamation
https://openreview.net/forum?id=VqzXzA9hjaX
AUTHORS: Tianshu Huang, Tianlong Chen, Sijia Liu, Shiyu Chang, Lisa Amini, Zhangyang Wang
HIGHLIGHT: We study a new problem named Optimizer Amalgamation: combining a pool of "teacher" optimizers into a single "student" optimizer which can have stronger problem-specific performance.

550, TITLE: Variational autoencoders in the presence of low-dimensional data: landscape and implicit bias
https://openreview.net/forum?id=y_op4lLLaWL
AUTHORS: Frederic Koehler, Viraj Mehta, Chenghui Zhou, Andrej Risteski
HIGHLIGHT: We analyze the landscape and implicit optimization bias of VAEs in the presence of low-dimensional data and the implications thereof.

551, TITLE: Pareto Policy Adaptation
https://openreview.net/forum?id=wfZGut6e09
AUTHORS: Panagiotis Kyriakis, Jyotirmoy Deshmukh, Paul Bogdan
HIGHLIGHT: We propose a policy gradient method for multi-objective reinforcement learning under unknown, linear preferences.

552, TITLE: Vector-quantized Image Modeling with Improved VQGAN
https://openreview.net/forum?id=pfNyExj7z2
AUTHORS: Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, Yonghui Wu
HIGHLIGHT: We propose the ViT-VQGAN and further explore a Vector-quantized Image Modeling (VIM) approach on both generative and discriminative tasks on images.

553, TITLE: Decentralized Learning for Overparameterized Problems: A Multi-Agent Kernel Approximation Approach
https://openreview.net/forum?id=oj2yn1Q4Ett
AUTHORS: Prashant Khanduri, Haibo Yang, Mingyi Hong, Jia Liu, Hoi To Wai, Sijia Liu
HIGHLIGHT: Decentralized optimization for overparameterized kernel learning

554, TITLE: Pre-training Molecular Graph Representation with 3D Geometry
https://openreview.net/forum?id=xQUe1pOKPam
AUTHORS: Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, Jian Tang
HIGHLIGHT: We proposed a new SSL framework to make 3D geomety information helpful for 2D representation, in terms of the downstream tasks with 2D info only.

555, TITLE: GeneDisco: A Benchmark for Experimental Design in Drug Discovery
https://openreview.net/forum?id=-w2oomO6qgc
AUTHORS: Arash Mehrjou, Ashkan Soleymani, Andrew Jesson, Pascal Notin, Yarin Gal, Stefan Bauer, Patrick Schwab
HIGHLIGHT: A comprehensive benchmark for batch active learning in drug discovery.

556, TITLE: Phase Collapse in Neural Networks
https://openreview.net/forum?id=iPHLcmtietq
AUTHORS: Florentin Guth, John Zarka, St?phane Mallat
HIGHLIGHT: The classification accuracy of CNNs mostly relies on the mechanism of phase collapses to eliminate spatial variability and linearly separate class means.

557, TITLE: Surrogate NAS Benchmarks: Going Beyond the Limited Search Spaces of Tabular NAS Benchmarks
https://openreview.net/forum?id=OnpFa95RVqs
AUTHORS: Arber Zela, Julien Niklas Siems, Lucas Zimmer, Jovita Lukasik, Margret Keuper, Frank Hutter
HIGHLIGHT: We present surrogate benchmarks for neural architecture search and a general methodology for constructing them.

558, TITLE: DictFormer: Tiny Transformer with Shared Dictionary
https://openreview.net/forum?id=GWQWAeE9EpB
AUTHORS: Qian Lou, Ting Hua, Yen-Chang Hsu, Yilin Shen, Hongxia Jin
HIGHLIGHT: We propose DictFormer with efficient shared dictionary to provide a compact, fast, and accurate transformer model.

559, TITLE: Reverse Engineering of Imperceptible Adversarial Image Perturbations
https://openreview.net/forum?id=gpp7cf0xdfN
AUTHORS: Yifan Gong, Yuguang Yao, Yize Li, Yimeng Zhang, Xiaoming Liu, Xue Lin, Sijia Liu
HIGHLIGHT: Reverse engineer adversarial image perturbations with a denoiser-based framework.

560, TITLE: An Explanation of In-context Learning as Implicit Bayesian Inference
https://openreview.net/forum?id=RdJVFCHjUMI
AUTHORS: Sang Michael Xie, Aditi Raghunathan, Percy Liang, Tengyu Ma
HIGHLIGHT: In-context learning emerges both theoretically and empirically when the pretraining distribution is a mixture distribution, resulting in the language model implicitly performing Bayesian inference in its forward pass.

561, TITLE: Learning Efficient Image Super-Resolution Networks via Structure-Regularized Pruning
https://openreview.net/forum?id=AjGC97Aofee
AUTHORS: Yulun Zhang, Huan Wang, Can Qin, Yun Fu
HIGHLIGHT: Learning efficient compressed models for bother lightweight and large image super-resolution networks

562, TITLE: Domain Adversarial Training: A Game Perspective

https://openreview.net/forum?id=AwgtcUAhBq
AUTHORS: David Acuna, Marc T Law, Guojun Zhang, Sanja Fidler
HIGHLIGHT: A novel perspective on domain-adversarial training that leads to more stable and performant optimizers.

563, TITLE: Tracking the risk of a deployed model and detecting harmful distribution shifts
https://openreview.net/forum?id=Ro_zAjZppv
AUTHORS: Aleksandr Podkopaev, Aaditya Ramdas
HIGHLIGHT: In this work, we design simple sequential tools for testing if the difference between source (training) and target (test) distributions leads to a significant drop in a risk function of interest, like accuracy or calibration.

564, TITLE: On the Pitfalls of Analyzing Individual Neurons in Language Models
https://openreview.net/forum?id=8uz0EWPQlMu
AUTHORS: Omer Antverg, Yonatan Belinkov
HIGHLIGHT: We analyze and compare methods to rank neurons in hidden representations according to their relevance to morphologic attributes, and show some of their weaknesses.

565, TITLE: Learning Invariant Representations on Multilingual Language Models for Unsupervised Cross-Lingual Transfer
https://openreview.net/forum?id=k7-s5HSSPE5
AUTHORS: Ruicheng Xian, Heng Ji, Han Zhao
HIGHLIGHT: We propose an importance-weighted domain adaptation method for unsupervised cross-lingual learning that is effective at correcting class prior shifts, a distributional property that negatively affects transfer performance.

566, TITLE: Provable Adaptation across Multiway Domains via Representation Learning
https://openreview.net/forum?id=gRCCdgpVZf
AUTHORS: Zhili Feng, Shaobo Han, Simon Shaolei Du
HIGHLIGHT: We propose a model which consists of a domain-invariant latent representation layer and a domain-specific linear prediction layer with a low-rank tensor structure.

567, TITLE: Permutation-Based SGD: Is Random Optimal?
https://openreview.net/forum?id=YiBa9HKTyXE
AUTHORS: Shashank Rajput, Kangwook Lee, Dimitris Papailiopoulos
HIGHLIGHT: We show that the question of whether random permutations are optimal for permutation-based SGD is nuanced, and depends on the family of functions one is trying to optimize.

568, TITLE: P-Adapters: Robustly Extracting Factual Information from Language Models with Diverse Prompts
https://openreview.net/forum?id=DhzIU48OcZh
AUTHORS: Benjamin Newman, Prafulla Kumar Choubey, Nazneen Rajani
HIGHLIGHT: We make prompting large language models for factual information more robust by introducing P-Adapter models.

569, TITLE: Optimal Representations for Covariate Shift
https://openreview.net/forum?id=Rf58LPCwJj0
AUTHORS: Yangjun Ruan, Yann Dubois, Chris J. Maddison
HIGHLIGHT: We give a simple variational objective whose optima are exactly the set of representations that are robust under covariate shift

570, TITLE: Imitation Learning by Reinforcement Learning
https://openreview.net/forum?id=1zwleytEpYx
AUTHORS: Kamil Ciosek
HIGHLIGHT: For deterministic experts, you can do imitation learning by calling an RL solver once, with a stationary reward signal.

571, TITLE: Evolutionary Diversity Optimization with Clustering-based Selection for Reinforcement Learning
https://openreview.net/forum?id=74x5BXs4bWD
AUTHORS: Yutong Wang, Ke Xue, Chao Qian
HIGHLIGHT: We propose EDO-CS, a new Evolutionary Diversity Optimization algorithm with Clustering-based Selection that can achieve a set of policies with both high quality and diversity efficiently.

572, TITLE: Cross-Domain Imitation Learning via Optimal Transport
https://openreview.net/forum?id=xP3cPq2hQC

AUTHORS: Arnaud Fickinger, Samuel Cohen, Stuart Russell, Brandon Amos
HIGHLIGHT: We study the use of Gromov-Wasserstein for cross-domain imitation learning
We also release the set of considered benchmark tasks and implementations of our methods in order to further drive research in this challenging area.

573, TITLE: DemoDICE: Offline Imitation Learning with Supplementary Imperfect Demonstrations
https://openreview.net/forum?id=BrPdX1bDZkQ
AUTHORS: Geon-Hyeong Kim, Seokin Seo, Jongmin Lee, Wonseok Jeon, HyeongJoo Hwang, Hongseok Yang, Kee-Eung Kim
HIGHLIGHT: Under this setting, we propose DemoDICE, which effectively utilizes imperfect demonstrations by matching the stationary distribution of a policy with experts' distribution while penalizing its deviation from the overall demonstrations.

574, TITLE: Sqrt(d) Dimension Dependence of Langevin Monte Carlo
https://openreview.net/forum?id=5-2mX9_U5i
AUTHORS: Ruilin Li, Hongyuan Zha, Molei Tao
HIGHLIGHT: The known dimension dependence of LMC is improved, under regularity assumptions, from d to sqrt(d), based on a refined mean square analysis framework.

575, TITLE: Variational oracle guiding for reinforcement learning
https://openreview.net/forum?id=pjqqxepwoMy
AUTHORS: Dongqi Han, Tadashi Kozuno, Xufang Luo, Zhao-Yun Chen, Kenji Doya, Yuqing Yang, Dongsheng Li
HIGHLIGHT: We propose a variational Bayes framework leveraging oracle (hindsight) information available in training to improve deep reinforcement learning.

576, TITLE: Graph Neural Networks with Learnable Structural and Positional Representations
https://openreview.net/forum?id=wTTjnvGphYj
AUTHORS: Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, Xavier Bresson
HIGHLIGHT: We propose a novel GNN architecture (LSPE) which decouples structural and positional representations to make easy for the network to learn the two essential properties.

577, TITLE: Target-Side Data Augmentation for Sequence Generation
https://openreview.net/forum?id=pz1euXohm4H
AUTHORS: Shufang Xie, Ang Lv, Yingce Xia, Lijun Wu, Tao Qin, Rui Yan, Tie-Yan Liu
HIGHLIGHT: We study the data augmentation for target-side conditional input of autoregressive sequence generation and propose a new method to build soft synthetic data.

578, TITLE: BadPre: Task-agnostic Backdoor Attacks to Pre-trained NLP Foundation Models
https://openreview.net/forum?id=Mng8CQ9eBW
AUTHORS: Kangjie Chen, Yuxian Meng, Xiaofei Sun, Shangwei Guo, Tianwei Zhang, Jiwei Li, Chun Fan
HIGHLIGHT: In this work, we propose BadPre, the first backdoor attack against various downstream models built based on pre-trained NLP models.

579, TITLE: Provably Robust Adversarial Examples
https://openreview.net/forum?id=UMfhoMtIaP5
AUTHORS: Dimitar Iliev Dimitrov, Gagandeep Singh, Timon Gehr, Martin Vechev
HIGHLIGHT: We introduce the concept of provably robust adversarial examples $-$ connected input regions constructed from standard adversarial examples and guaranteed to be provably robust to a set of perturbations.

580, TITLE: LORD: Lower-Dimensional Embedding of Log-Signature in Neural Rough Differential Equations
https://openreview.net/forum?id=fCG75wd39ze
AUTHORS: JAEHOON LEE, Jinsung Jeon, Sheo yon Jhin, Jihyeon Hyeong, Jayoung Kim, Minju Jo, Kook Seungji, Noseong Park
HIGHLIGHT: We reduce the complexity of processing higher depth log-signatures in NRDE.

581, TITLE: Ada-NETS: Face Clustering via Adaptive Neighbour Discovery in the Structure Space
https://openreview.net/forum?id=QJWVP4CTmW4
AUTHORS: Yaohua Wang, Yaobin Zhang, Fangyi Zhang, Senzhang Wang, Ming Lin, YuQi Zhang, Xiuyu Sun
HIGHLIGHT: A novel algorithm named Ada-NETS is proposed to construct the clean graph for GCNs to cluster faces in this paper.

582, TITLE: Prototypical Contrastive Predictive Coding
https://openreview.net/forum?id=8la28hZOwug
AUTHORS: Kyungmin Lee
HIGHLIGHT: We propose prototypical contrastive predictive coding for efficient distillation of representational knowledge of one network into other network.

583, TITLE: Neural Models for Output-Space Invariance in Combinatorial Problems
https://openreview.net/forum?id=ibrUkC-pbis
AUTHORS: Yatin Nandwani, Vidit Jain, Mausam ., Parag Singla
HIGHLIGHT: We design GNN based neural models to achieve output space invariance in combinatorial problems, e.g. solving 16 x 16 sudoku after training on 9 x 9 sudoku

584, TITLE: FILIP: Fine-grained Interactive Language-Image Pre-Training
https://openreview.net/forum?id=cpDhcsEDC2
AUTHORS: Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, Chunjing Xu
HIGHLIGHT: We introduce a large-scale Fine-grained Interacitve Language-Image Pretraining (FILIP) to achieve finer-level alignment through a new cross-modal late interaction mechanism, which can boost the performance on more grounded vision and language tasks.
Furthermore, we construct a new large-scale image-text pair dataset called FILIP300M for pre-training.

585, TITLE: Evaluating Model-Based Planning and Planner Amortization for Continuous Control
https://openreview.net/forum?id=SS8F6tFX3-
AUTHORS: Arunkumar Byravan, Leonard Hasenclever, Piotr Trochim, Mehdi Mirza, Alessandro Davide Ialongo, Yuval Tassa, Jost Tobias Springenberg, Abbas Abdolmaleki, Nicolas Heess, Josh Merel, Martin Riedmiller
HIGHLIGHT: We combine MPC with model-free RL and evaluate on continuous control tasks from scratch and in transfer settings; our results show that model-free RL is a strong baseline in single task settings and model-based methods shine in multi-goal tasks.

586, TITLE: Embedded-model flows: Combining the inductive biases of model-free deep learning and explicit probabilistic modeling
https://openreview.net/forum?id=9pEJSVfDbba
AUTHORS: Gianluigi Silvestri, Emily Fertig, Dave Moore, Luca Ambrogioni
HIGHLIGHT: We introduce bijective transformations that embed domain-specific inductive biases in Normalizing Flow architectures.

587, TITLE: On Redundancy and Diversity in Cell-based Neural Architecture Search
https://openreview.net/forum?id=rFJWoYoxrDB
AUTHORS: Xingchen Wan, Binxin Ru, Pedro M Esperan?a, Zhenguo Li
HIGHLIGHT: We analyse and explore the redundancies and diversity of popular cell-based search spaces in NAS.

588, TITLE: Hierarchical Variational Memory for Few-shot Learning Across Domains
https://openreview.net/forum?id=i3RI65sR7N
AUTHORS: Yingjun Du, Xiantong Zhen, Ling Shao, Cees G. M. Snoek
HIGHLIGHT: We propose a hierarchical memory that stores features at different semantic levels for the cross-domain few-shot learning.

589, TITLE: DARA: Dynamics-Aware Reward Augmentation in Offline Reinforcement Learning
https://openreview.net/forum?id=9SDQB3b68K
AUTHORS: Jinxin Liu, Zhang Hongyin, Donglin Wang
HIGHLIGHT: In this paper, we thus 1) formulate the offline dynamics adaptation by using (source) offline data collected from another dynamics to relax the requirement for the extensive (target) offline data, 2) characterize the dynamics shift problem in which prior offline methods does not scale well, and 3) derive a simple dynamics-aware reward augmentation (DARA) framework from both model-free and model-based offline settings.

590, TITLE: Hidden Parameter Recurrent State Space Models For Changing Dynamics Scenarios
https://openreview.net/forum?id=ds8yZOUsea
AUTHORS: Vaisakh Shaj, Dieter B?chler, Rohit Sonker, Philipp Becker, Gerhard Neumann
HIGHLIGHT: A new formalism for extending Recurrent State Space Models (RSSMs) to changing dynamics scenarios.

591, TITLE: Information Gain Propagation: a New Way to Graph Active Learning with Soft Labels

https://openreview.net/forum?id=USC0-nvGPK
AUTHORS: Wentao Zhang, Yexin Wang, Zhenbang You, Meng Cao, Ping Huang, Jiulong Shan, Zhi Yang, Bin CUI
HIGHLIGHT: A new way to graph active learning with soft labels

592, TITLE: GLASS: GNN with Labeling Tricks for Subgraph Representation Learning
https://openreview.net/forum?id=XLxhEjKNbXj
AUTHORS: Xiyuan Wang, Muhan Zhang
HIGHLIGHT: With this insight, we introduce an expressive and scalable labeling trick, namely max-zero-one, and propose GLASS (GNN with LAbel for SubgraphS).

593, TITLE: Causal Contextual Bandits with Targeted Interventions
https://openreview.net/forum?id=F5Em8ASCosV
AUTHORS: Chandrasekar Subramanian, Balaraman Ravindran
HIGHLIGHT: A new, more realistic, formalism of contextual bandits involving causal side-information and targeted interventions, along with a novel algorithm that exploits features of the new setting such as information leakage to learn good policies quickly.

594, TITLE: PF-GNN: Differentiable particle filtering based approximation of universal graph representations
https://openreview.net/forum?id=oh4TirnfSem
AUTHORS: Mohammed Haroon Dupty, Yanfei Dong, Wee Sun Lee
HIGHLIGHT: Increasing the expressive power of Graph Neural Networks by using techniques from exact isomorphism solvers with a particle filtering approach.

595, TITLE: Optimal ANN-SNN Conversion for High-accuracy and Ultra-low-latency Spiking Neural Networks
https://openreview.net/forum?id=7B3IJMM1k_M
AUTHORS: Tong Bu, Wei Fang, Jianhao Ding, PENGLIN DAI, Zhaofei Yu, Tiejun Huang
HIGHLIGHT: An ANN-SNN conversion method enables high-accuracy and ultra-low-latency deep SNNs.

596, TITLE: Switch to Generalize: Domain-Switch Learning for Cross-Domain Few-Shot Classification
https://openreview.net/forum?id=H-iABMvzIc
AUTHORS: Zhengdong Hu, Yifan Sun, Yi Yang
HIGHLIGHT: This paper considers few-shot learning under the cross-domain scenario.

597, TITLE: Denoising Likelihood Score Matching for Conditional Score-based Data Generation
https://openreview.net/forum?id=LcF-EEt8cCC
AUTHORS: Chen-Hao Chao, Wei-Fang Sun, Bo-Wun Cheng, Yi-Chen Lo, Chia-Che Chang, Yu-Lun Liu, Yu-Lin Chang, Chia-Ping Chen, Chun-Yi Lee
HIGHLIGHT: In this paper, we theoretically formulate a new training objective, called Denoising Likelihood Score Matching (DLSM) loss, for the classifier to match the gradients of the true log likelihood density.

598, TITLE: Offline Reinforcement Learning with Value-based Episodic Memory
https://openreview.net/forum?id=RCZqv9NXlZ
AUTHORS: Xiaoteng Ma, Yiqin Yang, Hao Hu, Jun Yang, Chongjie Zhang, Qianchuan Zhao, Bin Liang, Qihan Liu
HIGHLIGHT: We propose a new offline RL method which uses expectile value learning and memory-based planning.

599, TITLE: Contrastive Clustering to Mine Pseudo Parallel Data for Unsupervised Translation
https://openreview.net/forum?id=pN1JOdrSY9
AUTHORS: Xuan-Phi Nguyen, Hongyu Gong, Yun Tang, Changhan Wang, Philipp Koehn, Shafiq Joty
HIGHLIGHT: We propose a fine-tuning loss that enables pre-trained model's ability to mine pseudo-parallel data for fully unsupervised machine translation.

600, TITLE: Proof Artifact Co-Training for Theorem Proving with Language Models
https://openreview.net/forum?id=rpxJc9j04U
AUTHORS: Jesse Michael Han, Jason Rute, Yuhuai Wu, Edward Ayers, Stanislas Polu
HIGHLIGHT: Co-training on low-level proof tasks in a Transformer theorem prover improves success from 32% to 48% on Lean theorems.

601, TITLE: Learning Disentangled Representation by Exploiting Pretrained Generative Models: A Contrastive Learning View
https://openreview.net/forum?id=j-63FSNcO5a

AUTHORS: Xuanchi Ren, Tao Yang, Yuwang Wang, Wenjun Zeng
HIGHLIGHT: DisCo is a new contrastive learning framework to leverage pretrained generative models to jointly learn disentangled representation and discover disentangled directions in the latent space.

602, TITLE: Efficient Self-supervised Vision Transformers for Representation Learning
https://openreview.net/forum?id=fVu3o-YUGQK
AUTHORS: Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, Jianfeng Gao
HIGHLIGHT: Achieving SoTA ImageNet linear probe task with 10 times higher throughput, using the synergy of a multi-stage Transformer architecture and a non-contrastive region-matching pre-training task.

603, TITLE: Accelerated Policy Learning with Parallel Differentiable Simulation
https://openreview.net/forum?id=ZSKRQMvttc
AUTHORS: Jie Xu, Miles Macklin, Viktor Makoviychuk, Yashraj Narang, Animesh Garg, Fabio Ramos, Wojciech Matusik
HIGHLIGHT: We propose an efficient policy learning method leveraging the recent advance of differentiable simulation, and our method outperforms state-of-the-art algorithms in both sample efficiency and wall clock time on multiple challenging control tasks.

604, TITLE: Bridging Recommendation and Marketing via Recurrent Intensity Modeling
https://openreview.net/forum?id=TZeArecH2Nf
AUTHORS: Yifei Ma, Ge Liu, Anoop Deoras
HIGHLIGHT: We repurpose an item-recommender system to recommend users for item providers for the purpose of content promotion and diversity.

605, TITLE: The Unreasonable Effectiveness of Random Pruning: Return of the Most Naive Baseline for Sparse Training
https://openreview.net/forum?id=VBZJ_3tz-t
AUTHORS: Shiwei Liu, Tianlong Chen, Xiaohan Chen, Li Shen, Decebal Constantin Mocanu, Zhangyang Wang, Mykola Pechenizkiy
HIGHLIGHT: We revisit the most naive baseline, random pruning, in sparse training and highlight a perhaps counter-intuitive finding: random pruning can be quite powerful for the sparse training of modern neural networks.

606, TITLE: Sparsity Winning Twice: Better Robust Generalization from More Efficient Training
https://openreview.net/forum?id=SYuJXrXq8tw
AUTHORS: Tianlong Chen, Zhenyu Zhang, pengjun wang, Santosh Balachandra, Haoyu Ma, Zehao Wang, Zhangyang Wang
HIGHLIGHT: In this paper, we investigate this intriguing problem from a new perspective, i.e., $\textit{injecting appropriate forms of sparsity}$ during adversarial training.

607, TITLE: The Rich Get Richer: Disparate Impact of Semi-Supervised Learning
https://openreview.net/forum?id=DXPftn5kjQK
AUTHORS: Zhaowei Zhu, Tianyi Luo, Yang Liu
HIGHLIGHT: We reveal the disparate impacts of deploying SSL: the "rich" sub-population (higher baseline accuracy without SSL) benefits more from SSL; while the "poor" sub-population (low baseline accuracy) might even observe a performance drop after SSL.

608, TITLE: StyleNeRF: A Style-based 3D Aware Generator for High-resolution Image Synthesis
https://openreview.net/forum?id=iUuzzTMUw9K
AUTHORS: Jiatao Gu, Lingjie Liu, Peng Wang, Christian Theobalt
HIGHLIGHT: We present StyleNeRF, a 3D-aware generative model that can synthesize high-resolution images with high multi-view consistency.

609, TITLE: AEVA: Black-box Backdoor Detection Using Adversarial Extreme Value Analysis
https://openreview.net/forum?id=OM_lYiHXiCL
AUTHORS: Junfeng Guo, Ang Li, Cong Liu
HIGHLIGHT: Based on this observation, we propose the adversarial extreme value analysis(AEVA) algorithm to detect backdoors in black-box neural networks.

610, TITLE: Generalized Kernel Thinning
https://openreview.net/forum?id=IfNu7Dr-3fQ
AUTHORS: Raaz Dwivedi, Lester Mackey

HIGHLIGHT: By generalizing kernel thinning, we develop new tools for compressing distributions more effectively than i.i.d. sampling and establish new dimension-free, $O(\sqrt{\log n/n})$ improvements over the usual $n^{-1/4}$ Monte Carlo integration error.

611, TITLE: Spatial Graph Attention and Curiosity-driven Policy for Antiviral Drug Discovery
https://openreview.net/forum?id=kavTY__jxp
AUTHORS: Yulun Wu, Nicholas Choma, Andrew Deru Chen, Mikaela Cashman, Erica Teixeira Prates, Veronica G Melesse Vergara, Manesh B Shah, Austin Clyde, Thomas Brettin, Wibe Albert de Jong, Neeraj Kumar, Martha S Head, Rick L. Stevens, Peter Nugent, Daniel A Jacobson, James B Brown
HIGHLIGHT: We developed a reinforcement learning framework that advances in exploiting spatial and attributional molecular information as well as exploring novel and synthesizable chemical structures for the purpose of antiviral drug discovery.

612, TITLE: Missingness Bias in Model Debugging
https://openreview.net/forum?id=Te5ytkqsnl
AUTHORS: Saachi Jain, Hadi Salman, Eric Wong, Pengchuan Zhang, Vibhav Vineet, Sai Vemprala, Aleksander Madry
HIGHLIGHT: We investigate how current missingness approximations for model debugging can impose undesirable biases on the model predictions and hinder our ability to debug models, and we show how transformer-based architectures can side-step these issues.

613, TITLE: Unsupervised Discovery of Object Radiance Fields
https://openreview.net/forum?id=rwE8SshAlxw
AUTHORS: Hong-Xing Yu, Leonidas Guibas, Jiajun Wu
HIGHLIGHT: Inferring object-centric factorized 3D scene representations from a single image, learned without 3D geometry or segmentation supervision.

614, TITLE: On the Convergence of mSGD and AdaGrad for Stochastic Optimization
https://openreview.net/forum?id=g5tANwND04i
AUTHORS: ruinan Jin, Yu Xing, Xingkang He
HIGHLIGHT: A theoretical paper focusing on the investigation for the convergence of mSGD and AdaGrad optimization algorithms.

615, TITLE: Bayesian Framework for Gradient Leakage
https://openreview.net/forum?id=f2lrIbGx3x7
AUTHORS: Mislav Balunovic, Dimitar Iliev Dimitrov, Robin Staab, Martin Vechev
HIGHLIGHT: We propose a theoretical framework for analysis of the Bayes optimal adversary for gradient leakage, and perform evaluation of existing defenses.

616, TITLE: How many degrees of freedom do we need to train deep networks: a loss landscape perspective
https://openreview.net/forum?id=ChMLTGRjFcU
AUTHORS: Brett W Larsen, Stanislav Fort, Nic Becker, Surya Ganguli
HIGHLIGHT: A variety of recent works, spanning pruning, lottery tickets, and training within random subspaces, have shown that deep neural networks can be trained using far fewer degrees of freedom than the total number of parameters. We analyze this phenomenon for random subspaces by first examining the success probability of hitting a training loss sublevel set when training within a random subspace of a given training dimensionality.

617, TITLE: Optimizing Neural Networks with Gradient Lexicase Selection
https://openreview.net/forum?id=J_2xNmVcY4
AUTHORS: Li Ding, Lee Spector
HIGHLIGHT: We propose Gradient Lexicase Selection, an evolutionary optimization method that improves the generalization of deep neural networks.

618, TITLE: Generalization of Overparametrized Deep Neural Network Under Noisy Observations
https://openreview.net/forum?id=bZJbzaj_llP
AUTHORS: Namjoon Suh, Hyunouk Ko, Xiaoming Huo
HIGHLIGHT: Study the generalization of overparametrized deep neural network with relu activation function with noisy dataset.

619, TITLE: TPU-GAN: Learning temporal coherence from dynamic point cloud sequences
https://openreview.net/forum?id=FEBFJ98FKx
AUTHORS: Zijie Li, Tianqin Li, Amir Barati Farimani
HIGHLIGHT: We propose a GAN framework for super-resolution task on dynamic point cloud sequences.

620, TITLE: Modular Lifelong Reinforcement Learning via Neural Composition
https://openreview.net/forum?id=5XmLzdslFNN
AUTHORS: Jorge A Mendez, Harm van Seijen, ERIC EATON
HIGHLIGHT: We explore the problem of lifelong RL of functionally composable knowledge, and develop an algorithm that demonstrates zero-shot and forward transfer, avoidance of forgetting, and backward transfer in discrete 2-D and robotic manipulation domains.

621, TITLE: PER-ETD: A Polynomially Efficient Emphatic Temporal Difference Learning Method
https://openreview.net/forum?id=-HSOjDPfhBJ
AUTHORS: Ziwei Guan, Tengyu Xu, Yingbin Liang
HIGHLIGHT: We propose a new off-policy evaluation algorithm called PER-ETD (i.e., PEriodically Restarted Emphatic TD), which improves upon its precursor ETD with reduced variance and polynomial sample efficiency.

622, TITLE: Sparse Attention with Learning to Hash
https://openreview.net/forum?id=VGnOJhd5Q1q
AUTHORS: Zhiqing Sun, Yiming Yang, Shinjae Yoo
HIGHLIGHT: We propose a new strategy for sparse attention, namely LHA (Learning-to-Hash Attention), which directly learns separate parameterized hash functions for queries and keys, respectively.

623, TITLE: MetaMorph: Learning Universal Controllers with Transformers
https://openreview.net/forum?id=Opmqtk_GvYL
AUTHORS: Agrim Gupta, Linxi Fan, Surya Ganguli, Li Fei-Fei
HIGHLIGHT: We learn a transformer based general purpose controller for a modular robot design space which can zero-shot generalize to unseen variations in dynamics, kinematics, new morphologies and tasks.

624, TITLE: Orchestrated Value Mapping for Reinforcement Learning
https://openreview.net/forum?id=c87d0TS4yX
AUTHORS: Mehdi Fatemi, Arash Tavakoli
HIGHLIGHT: We present a general convergent class of RL algorithms based on combining arbitrary value mappings and reward decomposition.

625, TITLE: Efficient Sharpness-aware Minimization for Improved Training of Neural Networks
https://openreview.net/forum?id=n0OeTdNRG0Q
AUTHORS: Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, Vincent Tan
HIGHLIGHT: An efficient sharpness aware minimizer that improves the generalization

626, TITLE: An Operator Theoretic View On Pruning Deep Neural Networks
https://openreview.net/forum?id=pWBNOgdeURp
AUTHORS: William T Redman, MARIA FONOBEROVA, Ryan Mohr, Yannis Kevrekidis, Igor Mezic
HIGHLIGHT: Koopman operator theoretic methods are extended to deep neural network pruning and are shown to provide new insight into existing work.

627, TITLE: Mention Memory: incorporating textual knowledge into Transformers through entity mention attention
https://openreview.net/forum?id=OY1A8ejQgEX
AUTHORS: Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, William W. Cohen
HIGHLIGHT: Incorporate information from text corpus into Transformer model through within-model attention over table of entity mention representations.

628, TITLE: Maximum n-times Coverage for Vaccine Design
https://openreview.net/forum?id=ULfq0qR25dY
AUTHORS: Ge Liu, Alexander Dimitrakakis, Brandon Carter, David Gifford
HIGHLIGHT: We introduce the maximum n-times coverage problem that selects k overlays to maximize the summed coverage of weighted elements, where each element must be covered at least n times, and show its importance for vaccine design. We introduce the maximum n-times coverage problem that selects k overlays to maximize the summed coverage of weighted elements, where each element must be covered at least n times.

629, TITLE: Optimal Transport for Long-Tailed Recognition with Learnable Cost Matrix
https://openreview.net/forum?id=t98k9ePQQpn
AUTHORS: Hanyu Peng, Mingming Sun, Ping Li

HIGHLIGHT: A new paradigm for post-hoc correction based on optimal transport to cope with long-tailed recognition.

630, TITLE: Step-unrolled Denoising Autoencoders for Text Generation
https://openreview.net/forum?id=T0GpzBQ1Fg6
AUTHORS: Nikolay Savinov, Junyoung Chung, Mikolaj Binkowski, Erich Elsen, Aaron van den Oord
HIGHLIGHT: We propose a new generative model of text that unrolls the denoising process during training.

631, TITLE: Pseudo-Labeled Auto-Curriculum Learning for Semi-Supervised Keypoint Localization
https://openreview.net/forum?id=6Q52pZ-Th7N
AUTHORS: Can Wang, Sheng Jin, Yingda Guan, Wentao Liu, Chen Qian, Ping Luo, Wanli Ouyang
HIGHLIGHT: The first work that explores automatic curriculum learning for semi-supervised keypoint localization

632, TITLE: Surreal-GAN:Semi-Supervised Representation Learning via GAN for uncovering heterogeneous disease-related imaging patterns
https://openreview.net/forum?id=nf3A0WZsXS5
AUTHORS: Zhijian Yang, Junhao Wen, Christos Davatzikos
HIGHLIGHT: We proposed a novel method, Surreal-GAN, to derive low dimensional representation of disease-related patterns from neuroimaging data.

633, TITLE: NASI: Label- and Data-agnostic Neural Architecture Search at Initialization
https://openreview.net/forum?id=v-v1cpNNK_v
AUTHORS: Yao Shu, Shaofeng Cai, Zhongxiang Dai, Beng Chin Ooi, Bryan Kian Hsiang Low
HIGHLIGHT: To overcome this limitation, we propose a novel NAS algorithm called NAS at Initialization (NASI) that exploits the capability of a Neural Tangent Kernel in being able to characterize the performance of candidate architectures at initialization, hence allowing model training to be completely avoided to boost the search efficiency.

634, TITLE: A Relational Intervention Approach for Unsupervised Dynamics Generalization in Model-Based Reinforcement Learning
https://openreview.net/forum?id=YRq0ZUnzKoZ
AUTHORS: Jiaxian Guo, Mingming Gong, Dacheng Tao
HIGHLIGHT: This paper proposes a new model-based RL that could generalize to new environments.

635, TITLE: Structure-Aware Transformer Policy for Inhomogeneous Multi-Task Reinforcement Learning
https://openreview.net/forum?id=fy_XRVHqly
AUTHORS: Sunghoon Hong, Deunsol Yoon, Kee-Eung Kim
HIGHLIGHT: We present a modular Multi-task Reinforcement Learning method for inhomogeneous control tasks incorporating structural embedding of morphology.

636, TITLE: Generative Pseudo-Inverse Memory
https://openreview.net/forum?id=Harn4_EZBw
AUTHORS: Kha Pham, Hung Le, Man Ngo, Truyen Tran, Bao Ho, Svetha Venkatesh
HIGHLIGHT: We propose Generative Pseudo-Inverse Memory (GPM), a class of deep generative memory models that are fast to write in and read out.

637, TITLE: Optimization inspired Multi-Branch Equilibrium Models
https://openreview.net/forum?id=nbC8iTTXIrk
AUTHORS: Mingjie Li, Yisen Wang, Xingyu Xie, Zhouchen Lin
HIGHLIGHT: A Multi-Branch DEQ model and its training strategy proposed by minimizing an objective function designed as our purpose.

638, TITLE: Learning Versatile Neural Architectures by Propagating Network Codes
https://openreview.net/forum?id=KEQl-MZ5fg7
AUTHORS: Mingyu Ding, Yuqi Huo, Haoyu Lu, Linjie Yang, Zhe Wang, Zhiwu Lu, Jingdong Wang, Ping Luo
HIGHLIGHT: An efficient NAS method by inverting neural predictors to directly update architectures and a multitask NAS benchmark for cross-task architecture design and analysis.
We first introduce a unified design space for multiple tasks and build a multitask NAS benchmark (NAS-Bench-MR) on many wildly used datasets, including ImageNet, Cityscapes, KITTI, and HMDB51.

639, TITLE: Learning to Downsample for Segmentation of Ultra-High Resolution Images
https://openreview.net/forum?id=HndgQudNb91

AUTHORS: Chen Jin, Ryutaro Tanno, Thomy Mertzanidou, Eleftheria Panagiotaki, Daniel C. Alexander
HIGHLIGHT: We propose a method for learning to downsample ultra high-resolution images that reflects the importance of each location.

640, TITLE: AS-MLP: An Axial Shifted MLP Architecture for Vision
https://openreview.net/forum?id=fvLLcIYmXb
AUTHORS: Dongze Lian, Zehao Yu, Xing Sun, Shenghua Gao
HIGHLIGHT: We design the first MLP-based architecture for downstream tasks. It achieves competitive performance compared to the transformer-based architecture, which establishes a new strong baseline of MLP-based architecture.

641, TITLE: Unsupervised Disentanglement with Tensor Product Representations on the Torus
https://openreview.net/forum?id=neqU3HWDgE
AUTHORS: Michael Rotman, Amit Dekel, Shir Gur, Yaron Oz, Lior Wolf
HIGHLIGHT: Decomposition of a latent space on a torus leads to a disentangled representation

642, TITLE: Training Data Generating Networks: Shape Reconstruction via Bi-level Optimization
https://openreview.net/forum?id=dDo8druYppX
AUTHORS: Biao Zhang, Peter Wonka
HIGHLIGHT: We propose a novel 3d shape representation for 3d shape reconstruction from a single image.

643, TITLE: NeuPL: Neural Population Learning
https://openreview.net/forum?id=MIX3fJkl_1
AUTHORS: Siqi Liu, Luke Marris, Daniel Hennes, Josh Merel, Nicolas Heess, Thore Graepel
HIGHLIGHT: We propose NeuPL, a general and efficient population learning framework that learns and represents diverse policies in symmetric zero-sum games within a single conditional network via self-play.

644, TITLE: Delaunay Component Analysis for Evaluation of Data Representations
https://openreview.net/forum?id=HTVch9AMPa
AUTHORS: Petra Poklukar, Vladislav Polianskii, Anastasiia Varava, Florian T. Pokorny, Danica Kragic Jensfelt
HIGHLIGHT: We present Delaunay Component Analysis (DCA) framework for evaluation of learned data representations which anayzes geometric and topological properties of representation spaces using Delaunay graphs.

645, TITLE: Mapping conditional distributions for domain adaptation under generalized target shift
https://openreview.net/forum?id=sPfB2PI87BZ
AUTHORS: Matthieu Kirchmeyer, Alain Rakotomamonjy, Emmanuel de Bezenac, patrick gallinari
HIGHLIGHT: We propose a novel theoretically grounded approach for domain adaptation under Generalized Target Shift; it learns a map between pretrained source and target representations that matches conditional distributions and recovers target proportions.

646, TITLE: Visual Representation Learning over Latent Domains
https://openreview.net/forum?id=kG0AtPi6JI1
AUTHORS: Lucas Deecke, Timothy Hospedales, Hakan Bilen
HIGHLIGHT: A new setting for learning over data with multiple latent domains is introduced, alongside new methods for it.

647, TITLE: Scale Efficiently: Insights from Pretraining and Finetuning Transformers
https://openreview.net/forum?id=f2OYVDyfIB
AUTHORS: Yi Tay, Mostafa Dehghani, Jinfeng Rao, William Fedus, Samira Abnar, Hyung Won Chung, Sharan Narang, Dani Yogatama, Ashish Vaswani, Donald Metzler
HIGHLIGHT: Scaling laws for upstream and downstream tasks

648, TITLE: A Theory of Tournament Representations
https://openreview.net/forum?id=zzk231Ms1Ih
AUTHORS: Arun Rajkumar, Vishnu Veerathu, Abdul Bakey Mir
HIGHLIGHT: We develop a theory to understand tournament representations i.e. structurally characterise when a tournament graph can be represented in lower dimensions using a skew symmetric matrix.

649, TITLE: Self-Joint Supervised Learning
https://openreview.net/forum?id=zuqcmNVK4c2
AUTHORS: Navid Kardan, Mubarak Shah, Mitch Hill

HIGHLIGHT: To address these limitations, we propose a new supervised learning paradigm called self-joint learning that generalizes the standard approach by modeling the joint conditional distribution of two observed samples, where each sample is an image and its label.

650, TITLE: A Unified Contrastive Energy-based Model for Understanding the Generative Ability of Adversarial Training
https://openreview.net/forum?id=XhF2VOMRHS
AUTHORS: Yifei Wang, Yisen Wang, Jiansheng Yang, Zhouchen Lin
HIGHLIGHT: In this paper, we demystify this phenomenon by developing a unified probabilistic framework, called Contrastive Energy-based Models (CEM).

651, TITLE: Bayesian Neural Network Priors Revisited
https://openreview.net/forum?id=xkjqJYqRJy
AUTHORS: Vincent Fortuin, Adri? Garriga-Alonso, Sebastian W. Ober, Florian Wenzel, Gunnar Ratsch, Richard E Turner, Mark van der Wilk, Laurence Aitchison
HIGHLIGHT: Using BNN priors that are not isotropic Gaussians can improve performance and reduce the cold posterior effect.

652, TITLE: On Non-Random Missing Labels in Semi-Supervised Learning
https://openreview.net/forum?id=6yVvwR9H9Oj
AUTHORS: Xinting Hu, Yulei Niu, Chunyan Miao, Xian-Sheng Hua, Hanwang Zhang
HIGHLIGHT: We presented a principled class-aware doubly robust solution to handle the non-random missing labels in semi-supervised learning.

653, TITLE: Reward Uncertainty for Exploration in Preference-based Reinforcement Learning
https://openreview.net/forum?id=OWZVD-l-ZrC
AUTHORS: Xinran Liang, Katherine Shu, Kimin Lee, Pieter Abbeel
HIGHLIGHT: We present an exploration method specifically for preference-based RL algorithms.

654, TITLE: Learning to Remember Patterns: Pattern Matching Memory Networks for Traffic Forecasting
https://openreview.net/forum?id=wwDg3bbYBIq
AUTHORS: Hyunwook Lee, Seungmin Jin, Hyeshin Chu, Hongkyu Lim, Sungahn Ko
HIGHLIGHT: We presents a novel graph convolutional memory network, PM-MemNet, for the traffic forecasting.

655, TITLE: How to Inject Backdoors with Better Consistency: Logit Anchoring on Clean Data
https://openreview.net/forum?id=Bn09TnDngN
AUTHORS: Zhiyuan Zhang, Lingjuan Lyu, Weiqiang Wang, Lichao Sun, Xu Sun
HIGHLIGHT: We propose a novel logit anchoring approach for better global and instance-wise consistency in backdoor learning.

656, TITLE: Unified Visual Transformer Compression
https://openreview.net/forum?id=9jsZiUgkCZP
AUTHORS: Shixing Yu, Tianlong Chen, Jiayi Shen, Huan Yuan, Jianchao Tan, Sen Yang, Ji Liu, Zhangyang Wang
HIGHLIGHT: This paper proposes a unified ViT compression framework that seamlessly assembles three effective techniques: pruning, layer skipping, and knowledge distillation, which outperforms existing competitors.

657, TITLE: Pretraining Text Encoders with Adversarial Mixture of Training Signal Generators
https://openreview.net/forum?id=sX3XaHwotOg
AUTHORS: Yu Meng, Chenyan Xiong, Payal Bajaj, saurabh tiwary, Paul N. Bennett, Jiawei Han, Xia Song
HIGHLIGHT: We present AMOS, a new method that pretrains text encoders with an Adversarial learning curriculum via a Mixture Of Signals from multiple auxiliary generators.

658, TITLE: Knowledge Infused Decoding
https://openreview.net/forum?id=upnDJ7itech
AUTHORS: Ruibo Liu, Guoqing Zheng, Shashank Gupta, Radhika Gaonkar, Chongyang Gao, Soroush Vosoughi, Milad Shokouhi, Ahmed Hassan Awadallah
HIGHLIGHT: We propose a new decoding algorithm for language model generation, to obtain better performance in knowledge-intensive tasks.

659, TITLE: A Zest of LIME: Towards Architecture-Independent Model Distances
https://openreview.net/forum?id=OUz_9TiTv9j

AUTHORS: Hengrui Jia, Hongyu Chen, Jonas Guan, Ali Shahin Shamsabadi, Nicolas Papernot
HIGHLIGHT: We propose an architecture-independent distance metric that measures the similarity between ML models by comparing their global behaviors, approximated using LIME.

660, TITLE: Active Hierarchical Exploration with Stable Subgoal Representation Learning
https://openreview.net/forum?id=sNuFKTMktcY
AUTHORS: Siyuan Li, Jin Zhang, Jianhao Wang, Yang Yu, Chongjie Zhang
HIGHLIGHT: We propose a regularization to stabilize subgoal representation learning in goal-conditioned HRL and develop an active exploration strategy upon this stable representation.

661, TITLE: ViDT: An Efficient and Effective Fully Transformer-based Object Detector
https://openreview.net/forum?id=w4cXZDDib1H
AUTHORS: Hwanjun Song, Deqing Sun, Sanghyuk Chun, Varun Jampani, Dongyoon Han, Byeongho Heo, Wonjae Kim, Ming-Hsuan Yang
HIGHLIGHT: We integrate vision and detection transformers to build an efficient and effective fully transformer-based object detector.

662, TITLE: miniF2F: a cross-system benchmark for formal Olympiad-level mathematics
https://openreview.net/forum?id=9ZPegFuFTFv
AUTHORS: Kunhao Zheng, Jesse Michael Han, Stanislas Polu
HIGHLIGHT: We present $\textsf{miniF2F}$, a dataset of formal Olympiad-level mathematics problems statements intended to provide a unified cross-system benchmark for neural theorem proving.

663, TITLE: Neural Solvers for Fast and Accurate Numerical Optimal Control
https://openreview.net/forum?id=m8bypnj7Yl5
AUTHORS: Federico Berto, Stefano Massaroli, Michael Poli, Jinkyoo Park
HIGHLIGHT: We propose a hypersolvers approach for numerical optimal control which shows consistent Pareto improvements in solution accuracy and control performance.

664, TITLE: Evaluating Language-biased image classification based on semantic compositionality
https://openreview.net/forum?id=xNO7OEIcJc6
AUTHORS: Yoann Lemesle, Masataka Sawayama, Guillermo Valle-Perez, Maxime Adolphe, H?l?ne Sauz?on, Pierre-Yves Oudeyer
HIGHLIGHT: We developed a benchmark test based on hierarchical semantic compositionality to evaluate the language-biased image classification of artificial models and evaluated the CLIP model using it.

665, TITLE: Dealing with Non-Stationarity in MARL via Trust-Region Decomposition
https://openreview.net/forum?id=XHUxf5aRB3s
AUTHORS: Wenhao Li, Xiangfeng Wang, Bo Jin, Junjie Sheng, Hongyuan Zha
HIGHLIGHT: In this paper, we introduce a novel notion, the - measurement, to explicitly measure the non-stationarity of a policy sequence, which can be further proved to be bounded by the KL-divergence of consecutive joint policies.

666, TITLE: Information Prioritization through Empowerment in Visual Model-based RL
https://openreview.net/forum?id=DfUjyyRW90
AUTHORS: Homanga Bharadhwaj, Mohammad Babaeizadeh, Dumitru Erhan, Sergey Levine
HIGHLIGHT: Empowerment along with mutual information maximization helps learn functionally relevant factors in visual model-based RL

667, TITLE: Differentially Private Fractional Frequency Moments Estimation with Polylogarithmic Space
https://openreview.net/forum?id=7I8LPkcx8V
AUTHORS: Lun Wang, Iosif Pinelis, Dawn Song
HIGHLIGHT: We prove that $\mathbb{F}_p$ sketch, a well-celebrated streaming algorithm for frequency moments estimation, is differentially private as is when $p\in(0, 1]$.

668, TITLE: Learning Curves for SGD on Structured Features
https://openreview.net/forum?id=WPI2vbkAl3Q
AUTHORS: Blake Bordelon, Cengiz Pehlevan
HIGHLIGHT: The average case test risk for stochastic gradient descent on mean square error is computed in terms of feature covariance structure.

669, TITLE: What Makes Better Augmentation Strategies? Augment Difficult but Not too Different
https://openreview.net/forum?id=Ucx3DQbC9GH
AUTHORS: Jaehyung Kim, Dongyeop Kang, Sungsoo Ahn, Jinwoo Shin
HIGHLIGHT: Effective learning-based augmentation in NLP tasks by constructing difficult but not too different samples

670, TITLE: Multi-Mode Deep Matrix and Tensor Factorization
https://openreview.net/forum?id=6YVIk0sAkF_
AUTHORS: Jicong Fan
HIGHLIGHT: This paper presents a framework of multi-mode deep matrix and tensor factorizations to explore and exploit the full nonlinearity of the data in matrices and tensors.

671, TITLE: A Theoretical Analysis on Feature Learning in Neural Networks: Emergence from Inputs and Advantage over Fixed Features
https://openreview.net/forum?id=wMpS-Z_AI_E
AUTHORS: Zhenmei Shi, Junyi Wei, Yingyu Liang
HIGHLIGHT: Theoretical analysis of feature learning in neural networks on practically motivated learning problems, showing it strongly depends on the structure of the input distribution and leads to the advantage over fixed feature methods like kernels.

672, TITLE: A Unified Wasserstein Distributional Robustness Framework for Adversarial Training
https://openreview.net/forum?id=Dzpe9C1mpiv
AUTHORS: Anh Tuan Bui, Trung Le, Quan Hung Tran, He Zhao, Dinh Phung
HIGHLIGHT: A Unified Wasserstein Distributional Robustness Framework for Adversarial Training

673, TITLE: Measuring CLEVRness: Black-box Testing of Visual Reasoning Models
https://openreview.net/forum?id=UtGtoS4CYU
AUTHORS: Spyridon Mouselinos, Henryk Michalewski, Mateusz Malinowski
HIGHLIGHT: Black box testing of visual reasoning models as a two-player game.

674, TITLE: Latent Image Animator: Learning to animate image via latent space navigation
https://openreview.net/forum?id=7r6kDq0mK_
AUTHORS: Yaohui Wang, Di Yang, Francois Bremond, Antitza Dantcheva
HIGHLIGHT: Image animation via latent space navigation

675, TITLE: Learning Value Functions from Undirected State-only Experience
https://openreview.net/forum?id=6Pe99Juo9gd
AUTHORS: Matthew Chang, Arjun Gupta, Saurabh Gupta
HIGHLIGHT: We present a method for offline value learning without action labels.

676, TITLE: PI3NN: Out-of-distribution-aware Prediction Intervals from Three Neural Networks
https://openreview.net/forum?id=NoB8YgRuoFU
AUTHORS: Siyan Liu, Pei Zhang, Dan Lu, Guannan Zhang
HIGHLIGHT: We propose a novel prediction interval (PI) method for uncertainty quantification, which addresses three major issues with the state-of-the-art PI methods.

677, TITLE: The Effects of Invertibility on the Representational Complexity of Encoders in Variational Autoencoders
https://openreview.net/forum?id=7_JR7WpwKV1
AUTHORS: Divyansh Pareek, Andrej Risteski
HIGHLIGHT: VAEs with invertible mean map have small approximate encoders; non-invertible maps can result in large encoders.

678, TITLE: Bundle Networks: Fiber Bundles, Local Trivializations, and a Generative Approach to Exploring Many-to-one Maps
https://openreview.net/forum?id=aBXzcPPOuX
AUTHORS: Nico Courts, Henry Kvinge
HIGHLIGHT: We draw from the theory of fiber bundles in (differential) topology to create a principled approach to generative models that allow us to learn and sample from the "fiber" over a point in a many-to-one map.

679, TITLE: Direct then Diffuse: Incremental Unsupervised Skill Discovery for State Covering and Goal Reaching
https://openreview.net/forum?id=25kzAhUB1lz
AUTHORS: Pierre-Alexandre Kamienny, Jean Tarbouriech, Alessandro Lazaric, Ludovic Denoyer

HIGHLIGHT: In this paper, we build on the mutual information framework for skill discovery and introduce UPSIDE, which addresses the coverage-directedness trade-off in the following ways: 1) We design policies with a decoupled structure of a directed skill, trained to reach a specific region, followed by a diffusing part that induces a local coverage.

680, TITLE: A Biologically Interpretable Graph Convolutional Network to Link Genetic Risk Pathways and Imaging Phenotypes of Disease
https://openreview.net/forum?id=Lwr8We4MIxn
AUTHORS: Sayan Ghosal, Qiang Chen, Giulio Pergola, Aaron L Goldman, William Ulrich, Daniel R Weinberger, Archana Venkataraman
HIGHLIGHT: Biologically Informed Imaging-Genetics

681, TITLE: FILM: Following Instructions in Language with Modular Methods
https://openreview.net/forum?id=qI4542Y2s1D
AUTHORS: So Yeon Min, Devendra Singh Chaplot, Pradeep Kumar Ravikumar, Yonatan Bisk, Ruslan Salakhutdinov
HIGHLIGHT: We propose a modular method for embodied instruction following; our method achieves SOTA on the ALFRED benchmark by a large margin while using less data by eschewing both expert trajectories and low-level instructions.

682, TITLE: Stochastic Training is Not Necessary for Generalization
https://openreview.net/forum?id=ZBESeIUB5k
AUTHORS: Jonas Geiping, Micah Goldblum, Phil Pope, Michael Moeller, Tom Goldstein
HIGHLIGHT: Models trained with full-batch gradient descent and explicit regularization can match the generalization performance of models trained with stochastic minibatching.

683, TITLE: Multi-objective Optimization by Learning Space Partition
https://openreview.net/forum?id=FlwzVjfMryn
AUTHORS: Yiyang Zhao, Linnan Wang, Kevin Yang, Tianjun Zhang, Tian Guo, Yuandong Tian
HIGHLIGHT: Multi-objective Optimization by Learning Space Partition

684, TITLE: Sample Efficient Stochastic Policy Extragradient Algorithm for Zero-Sum Markov Game
https://openreview.net/forum?id=IvepFxYRDG
AUTHORS: Ziyi Chen, Shaocong Ma, Yi Zhou
HIGHLIGHT: This paper proposes a fully decentralized, model-free, provably convergent, sample efficient stochastic policy extragradient algorithm with symmetric and private policy updates

685, TITLE: Practical Integration via Separable Bijective Networks
https://openreview.net/forum?id=NlObxR0rosG
AUTHORS: Christopher M Bender, Patrick Emmanuel, Michael K. Reiter, Junier Oliva
HIGHLIGHT: We explore a method that enables learning over hypervolumes within the data space.

686, TITLE: Spherical Message Passing for 3D Molecular Graphs
https://openreview.net/forum?id=givsRXsOt9r
AUTHORS: Yi Liu, Limei Wang, Meng Liu, Yuchao Lin, Xuan Zhang, Bora Oztekin, Shuiwang Ji
HIGHLIGHT: In this work, we conduct analyses in the spherical coordinate system (SCS) for the complete identification of 3D graph structures.

687, TITLE: Axiomatic Explanations for Visual Search, Retrieval, and Similarity Learning
https://openreview.net/forum?id=TqNsv1TuCX9
AUTHORS: Mark Hamilton, Scott Lundberg, Stephanie Fu, Lei Zhang, William T. Freeman
HIGHLIGHT: We show that cooperative game theory provides an axiomatic characterization of model interpretability for visual search, retrieval, and similarity learning architectures

688, TITLE: A fast and accurate splitting method for optimal transport: analysis and implementation
https://openreview.net/forum?id=fCSq8yrDkc
AUTHORS: Vien V. Mai, Jacob Lindb?ck, Mikael Johansson
HIGHLIGHT: We develop a fast and reliable method for solving large-scale optimal transport (OT) problems at an unprecedented combination of speed and accuracy.

689, TITLE: Benchmarking the Spectrum of Agent Capabilities
https://openreview.net/forum?id=1W0z96MFEoH
AUTHORS: Danijar Hafner

HIGHLIGHT: We introduce Crafter, an open world survival game with visual inputs that evaluates a wide range of general abilities within a single environment.

690, TITLE: Noisy Feature Mixup
https://openreview.net/forum?id=vJb4I2ANmy
AUTHORS: Soon Hoe Lim, N. Benjamin Erichson, Francisco Utrera, Winnie Xu, Michael W. Mahoney
HIGHLIGHT: We propose and study Noisy Feature Mixup, a simple yet effective data augmentation method that leads to improved model robustness when compared to training with manifold mixup or noise injection alone.

691, TITLE: Learned Simulators for Turbulence
https://openreview.net/forum?id=msRBojTz-Nh
AUTHORS: Kim Stachenfeld, Drummond Buschman Fielding, Dmitrii Kochkov, Miles Cranmer, Tobias Pfaff, Jonathan Godwin, Can Cui, Shirley Ho, Peter Battaglia, Alvaro Sanchez-Gonzalez
HIGHLIGHT: Learned simulators that outperform baselines in capturing turbulent dynamics at low resolution across multiple challenging turbulence domains.

692, TITLE: Policy Gradients Incorporating the Future
https://openreview.net/forum?id=EHaUTlm2eHg
AUTHORS: David Venuto, Elaine Lau, Doina Precup, Ofir Nachum
HIGHLIGHT: Namely, we propose to allow an agent, during its training on past experience, to observe what \emph{actually} happened in the future at that time, while enforcing an information bottleneck to avoid the agent overly relying on this privileged information.

693, TITLE: Top-label calibration and multiclass-to-binary reductions
https://openreview.net/forum?id=WqoBaaPHS-
AUTHORS: Chirag Gupta, Aaditya Ramdas
HIGHLIGHT: We propose top-label calibration, a new and arguably natural notion for multiclass calibration, along with 'wrapper' calibration algorithms that reduce multiclass calibration to binary calibration.

694, TITLE: Neural Deep Equilibrium Solvers
https://openreview.net/forum?id=B0oHOwT5ENL
AUTHORS: Shaojie Bai, Vladlen Koltun, J Zico Kolter
HIGHLIGHT: A custom and lightweight neural solver for deep equilibrium models significantly improves their efficiency with minimal training.

695, TITLE: Conditional Contrastive Learning with Kernel
https://openreview.net/forum?id=AAJLBoGt0XM
AUTHORS: Yao-Hung Hubert Tsai, Tianqin Li, Martin Q. Ma, Han Zhao, Kun Zhang, Louis-Philippe Morency, Ruslan Salakhutdinov
HIGHLIGHT: This paper presents Conditional Contrastive Learning with Kernel (CCL-K) for conditional contrastive learning tasks under the scenario when we have insufficient data for some values of the condioning variable.

696, TITLE: Stability Regularization for Discrete Representation Learning
https://openreview.net/forum?id=6tmjoym9LR6
AUTHORS: Adeel Pervez, Efstratios Gavves
HIGHLIGHT: We present a method for training neural network models with discrete stochastic variables.

697, TITLE: Offline Reinforcement Learning with In-sample Q-Learning
https://openreview.net/forum?id=68n2s9ZJWF8
AUTHORS: Ilya Kostrikov, Ashvin Nair, Sergey Levine
HIGHLIGHT: Offline RL method with only dataset actions.

698, TITLE: It Takes Four to Tango: Multiagent Self Play for Automatic Curriculum Generation
https://openreview.net/forum?id=q4tZR1Y-UIs
AUTHORS: Yuqing Du, Pieter Abbeel, Aditya Grover
HIGHLIGHT: We propose Curriculum Self Play (CuSP), an automated goal generation framework that seeks to satisfy these desiderata by virtue of a multi-player game with 4 agents.

699, TITLE: Consistent Counterfactuals for Deep Models
https://openreview.net/forum?id=St6eyiTEHnG

AUTHORS: Emily Black, Zifan Wang, Matt Fredrikson
HIGHLIGHT: Counterfactual explanations are often inconsistent between virtually identical deep models. We introduce a new method to increase consistency while keeping costs low relative to other fixes.

700, TITLE: GradSign: Model Performance Inference with Theoretical Insights
https://openreview.net/forum?id=HObMhrCeAAF
AUTHORS: Zhihao Zhang, Zhihao Jia
HIGHLIGHT: Model performance inference inspired by sample-wise optimization landscape analysis

701, TITLE: Constructing Orthogonal Convolutions in an Explicit Manner
https://openreview.net/forum?id=Zr5W2LSRhD
AUTHORS: Tan Yu, Jun Li, YUNFENG CAI, Ping Li
HIGHLIGHT: In this work, we exploit the relation between the singular values of the convolution layer's Jacobian and the structure of the convolution kernel.

702, TITLE: Learning a subspace of policies for online adaptation in Reinforcement Learning
https://openreview.net/forum?id=4Muj-t_4o4
AUTHORS: Jean-Baptiste Gaya, Laure Soulier, Ludovic Denoyer
HIGHLIGHT: We propose an approach to learn a subspace of policies that are robust to different variations of the train environment.

703, TITLE: Towards Evaluating the Robustness of Neural Networks Learned by Transduction
https://openreview.net/forum?id=_5js_8uTrx1
AUTHORS: Jiefeng Chen, Xi Wu, Yang Guo, Yingyu Liang, Somesh Jha
HIGHLIGHT: Exploring evaluating the adversarial robustness of transductive-learning based defenses.

704, TITLE: On the Limitations of Multimodal VAEs
https://openreview.net/forum?id=w-CPUXXrAj
AUTHORS: Imant Daunhawer, Thomas M. Sutter, Kieran Chin-Cheong, Emanuele Palumbo, Julia E Vogt
HIGHLIGHT: We prove that the sub-sampling of modalities enforces an undesirable upper bound on the multimodal ELBO and thereby limits the generative quality of the respective models.

705, TITLE: Neural Parameter Allocation Search
https://openreview.net/forum?id=srtIXtySfT4
AUTHORS: Bryan A. Plummer, Nikoli Dryden, Julius Frost, Torsten Hoefler, Kate Saenko
HIGHLIGHT: An efficient approach for searching for the optimal allocation of parameters to layers across any neural network

706, TITLE: Distributionally Robust Models with Parametric Likelihood Ratios
https://openreview.net/forum?id=a34GrNaYEcS
AUTHORS: Paul Michel, Tatsunori Hashimoto, Graham Neubig
HIGHLIGHT: We learn adversarial parametric reweightings of the training data to reliably train more robust models

707, TITLE: Conditioning Sequence-to-sequence Networks with Learned Activations
https://openreview.net/forum?id=t5s-hd1bqLk
AUTHORS: Alberto Gil Couto Pimentel Ramos, Abhinav Mehrotra, Nicholas Donald Lane, Sourav Bhattacharya
HIGHLIGHT: Conditioning neural networks by learning the layer activations based on the conditioning vector

708, TITLE: Self-ensemble Adversarial Training for Improved Robustness
https://openreview.net/forum?id=oU3aTsmeRQV
AUTHORS: Hongjun Wang, Yisen Wang
HIGHLIGHT: This paper proposes an efficient self-ensemble method for adversarial trained classifiers and significantly improve their adversarial robustness

709, TITLE: X-model: Improving Data Efficiency in Deep Learning with A Minimax Model
https://openreview.net/forum?id=P3Bh01hBYTH
AUTHORS: Ximei Wang, Xinyang Chen, Jianmin Wang, Mingsheng Long
HIGHLIGHT: This paper proposes a novel X-Model for improving data efficiency in deep learning with a minimax model.

710, TITLE: Learning by Directional Gradient Descent

https://openreview.net/forum?id=5i7lJLuhTm
AUTHORS: David Silver, Anirudh Goyal, Ivo Danihelka, Matteo Hessel, Hado van Hasselt
HIGHLIGHT: Computing directional derivative of a recurrent function along a candidate direction, and using it to create a valid descent direction.

711, TITLE: SUMNAS: Supernet with Unbiased Meta-Features for Neural Architecture Search
https://openreview.net/forum?id=Z8FzvVU6_Kj
AUTHORS: Hyeonmin Ha, Ji-Hoon Kim, Semin Park, Byung-Gon Chun
HIGHLIGHT: We propose a supernet learning strategy that learns unbiased meta-features to tackle multi-model forgetting problem of neural architecture search.

712, TITLE: Low-Budget Active Learning via Wasserstein Distance: An Integer Programming Approach
https://openreview.net/forum?id=v8OlxjGn23S
AUTHORS: Rafid Mahmood, Sanja Fidler, Marc T Law
HIGHLIGHT: We propose an integer optimization problem for active learning and demonstrate how optimally selecting which points to label can significantly improve classifiers under low labeling budgets.

713, TITLE: Almost Tight L0-norm Certified Robustness of Top-k Predictions against Adversarial Perturbations
https://openreview.net/forum?id=gJLEXy3ySpu
AUTHORS: Jinyuan Jia, Binghui Wang, Xiaoyu Cao, Hongbin Liu, Neil Zhenqiang Gong
HIGHLIGHT: In this work, we derive the certified robustness against ℓ_0-norm adversarial perturbation for top-k prediction.

714, TITLE: SURF: Semi-supervised Reward Learning with Data Augmentation for Feedback-efficient Preference-based Reinforcement Learning
https://openreview.net/forum?id=TfhfZLQ2EJO
AUTHORS: Jongjin Park, Younggyo Seo, Jinwoo Shin, Honglak Lee, Pieter Abbeel, Kimin Lee
HIGHLIGHT: We present SURF, a semi-supervised reward learning algorithm with data augmentation for feedback-efficient preference-based RL.

715, TITLE: Rethinking Adversarial Transferability from a Data Distribution Perspective
https://openreview.net/forum?id=gVRhIEajG1k
AUTHORS: Yao Zhu, Jiacheng Sun, Zhenguo Li
HIGHLIGHT: In this paper, we rethink adversarial transferability from a data distribution perspective and further enhance transferability by score matching based optimization.

716, TITLE: Neural Networks as Kernel Learners: The Silent Alignment Effect
https://openreview.net/forum?id=1NvflqAdoom
AUTHORS: Alexander Atanasov, Blake Bordelon, Cengiz Pehlevan
HIGHLIGHT: Neural networks with small initialization, trained in the rich feature learning regime, can learn kernel regression solutions for a data adaptive kernel.

717, TITLE: Adaptive Control Flow in Transformers Improves Systematic Generalization
https://openreview.net/forum?id=KBQP4A_J1K
AUTHORS: R?bert Csord?s, Kazuki Irie, J?rgen Schmidhuber
HIGHLIGHT: We improve systematic generalization of Transformers on algorithmic tasks by introducing a novel attention mechanism and gating.

718, TITLE: Data Poisoning Won?t Save You From Facial Recognition
https://openreview.net/forum?id=B5XahNLmna
AUTHORS: Evani Radiya-Dixit, Sanghyun Hong, Nicholas Carlini, Florian Tramer
HIGHLIGHT: Data poisoning and adversarial examples won't protect users from facial recognition

719, TITLE: Procedural generalization by planning with self-supervised world models
https://openreview.net/forum?id=FmBegXJToY
AUTHORS: Ankesh Anand, Jacob C Walker, Yazhe Li, Eszter V?rtes, Julian Schrittwieser, Sherjil Ozair, Theophane Weber, Jessica B Hamrick
HIGHLIGHT: We study generalization in model-based agents and find that they excel at procedural generalization, with planning, self-supervision and data-diversity combining to yield SoTA results on Procgen; however, task generalization is more challenging.

720, TITLE: Global Convergence of Multi-Agent Policy Gradient in Markov Potential Games
https://openreview.net/forum?id=gfwON7rAm4
AUTHORS: Stefanos Leonardos, Will Overman, Ioannis Panageas, Georgios Piliouras
HIGHLIGHT: Convergence of policy gradient in a class of MDPs called Markov Potential Games in which cooperation is desired.

721, TITLE: Online Ad Hoc Teamwork under Partial Observability
https://openreview.net/forum?id=18Ys0-PzyPI
AUTHORS: Pengjie Gu, Mengchen Zhao, Jianye Hao, Bo An
HIGHLIGHT: To overcome partial observability, we introduce an information-based regularizer to derive \emph{proxy} representations of the learned variables from local observations.

722, TITLE: Stein Latent Optimization for Generative Adversarial Networks
https://openreview.net/forum?id=2-mkiUs9Jx7
AUTHORS: Uiwon Hwang, Heeseung Kim, Dahuin Jung, Hyemi Jang, Hyungyu Lee, Sungroh Yoon
HIGHLIGHT: We propose a novel GAN method that performs unsupervised conditional generation robustly on real-world datasets with balanced or imbalanced attributes even in the absence of attribute information (e.g., the imbalance ratio)

723, TITLE: A Johnson-Lindenstrauss Framework for Randomly Initialized CNNs
https://openreview.net/forum?id=YX0lrvdPQc
AUTHORS: Ido Nachum, Jan Hazla, Michael Gastpar, Anatoly Khina
HIGHLIGHT: We study how the geometric representation of a dataset change after the application of each randomly initialized layer of a neural network.

724, TITLE: Hot-Refresh Model Upgrades with Regression-Free Compatible Training in Image Retrieval
https://openreview.net/forum?id=HTp-6yLGGX
AUTHORS: Binjie Zhang, Yixiao Ge, Yantao Shen, Yu Li, Chun Yuan, XUYUAN XU, Yexin Wang, Ying Shan
HIGHLIGHT: We for the first time study the model regression problem in hot-refresh model upgrades of image retrieval systems with compatible representation learning.

725, TITLE: Objects in Semantic Topology
https://openreview.net/forum?id=d5SCUJ5t1k
AUTHORS: Shuo Yang, Peize Sun, Yi Jiang, Xiaobo Xia, Ruiheng Zhang, Zehuan Yuan, Changhu Wang, Ping Luo, Min Xu
HIGHLIGHT: In this paper, we provide a unified perspective: Semantic Topology.

726, TITLE: Sparse DETR: Efficient End-to-End Object Detection with Learnable Sparsity
https://openreview.net/forum?id=RRGVCN8kjim
AUTHORS: Byungseok Roh, JaeWoong Shin, Wuhyun Shin, Saehoon Kim
HIGHLIGHT: Sparse DETR is an efficient end-to-end object detector that sparsifies encoder queries by using the learnable decoder attention map predictor. It achieves better performance than Deformable DETR even with only 10% encoder queries on the COCO dataset.

727, TITLE: Partial Wasserstein Adversarial Network for Non-rigid Point Set Registration
https://openreview.net/forum?id=2ggNjUisGyr
AUTHORS: Ziming Wang, Nan Xue, Ling Lei, Gui-Song Xia
HIGHLIGHT: We propose a method for large scale partial distribution matching problem, and apply it to non-rigid point set registration task.

728, TITLE: On Distributed Adaptive Optimization with Gradient Compression
https://openreview.net/forum?id=CI-xXX9dg9l
AUTHORS: Xiaoyun Li, Belhal Karimi, Ping Li
HIGHLIGHT: We study COMP-AMS, a distributed optimization framework based on gradient averaging and adaptive AMSGrad algorithm.

729, TITLE: Implicit Bias of Adversarial Training for Deep Neural Networks
https://openreview.net/forum?id=l8It-0lE5e7
AUTHORS: Bochen Lv, Zhanxing Zhu
HIGHLIGHT: We provide theoretical understandings of the implicit bias imposed by adversarial training for homogeneous deep neural networks without explicit regularization.

730, TITLE: Disentanglement Analysis with Partial Information Decomposition
https://openreview.net/forum?id=pETy-HVvGtt
AUTHORS: Seiya Tokui, Issei Sato
HIGHLIGHT: We establish a framework to analyze information sharing in a multivariate representation with Partial Information Decomposition and propose a new disentanglement metric.

731, TITLE: Reliable Adversarial Distillation with Unreliable Teachers
https://openreview.net/forum?id=u6TRGdzhfip
AUTHORS: Jianing Zhu, Jiangchao Yao, Bo Han, Jingfeng Zhang, Tongliang Liu, Gang Niu, Jingren Zhou, Jianliang Xu, Hongxia Yang
HIGHLIGHT: Therefore, in this paper, we propose reliable introspective adversarial distillation (IAD) where students partially instead of fully trust their teachers.

732, TITLE: Knowledge Removal in Sampling-based Bayesian Inference
https://openreview.net/forum?id=dTqOcTUOQO
AUTHORS: Shaopeng Fu, Fengxiang He, Dacheng Tao
HIGHLIGHT: This paper proposes the first machine unlearning algorithm for MCMC.

733, TITLE: Entroformer: A Transformer-based Entropy Model for Learned Image Compression
https://openreview.net/forum?id=VrjOFfcnSV8
AUTHORS: Yichen Qian, Xiuyu Sun, Ming Lin, Zhiyu Tan, Rong Jin
HIGHLIGHT: In this work, we propose a novel transformer-based entropy model, termed Entroformer, to capture long-range dependencies in probability distribution estimation effectively and efficiently.

734, TITLE: Efficient Neural Causal Discovery without Acyclicity Constraints
https://openreview.net/forum?id=eYciPrLuUhG
AUTHORS: Phillip Lippe, Taco Cohen, Efstratios Gavves
HIGHLIGHT: We present ENCO, an efficient structure learning method that leverages observational and interventional data and scales to graphs with a thousand variables.

735, TITLE: Bandit Learning with Joint Effect of Incentivized Sampling, Delayed Sampling Feedback, and Self-Reinforcing User Preferences
https://openreview.net/forum?id=Q83vFlie_Pr
AUTHORS: Tianchen Zhou, Jia Liu, Chaosheng Dong, Yi Sun
HIGHLIGHT: In this paper, we consider a new multi-armed bandit (MAB) framework motivated by three common complications in online recommender systems in practice: (i) the platform (learning agent) cannot sample an intended product directly and has to incentivize customers to select this product (e.g., promotions and coupons); (ii) customer feedbacks are often received later than their selection times; and (iii) customer preferences among products are influenced and reinforced by historical feedbacks.

736, TITLE: A Conditional Point Diffusion-Refinement Paradigm for 3D Point Cloud Completion
https://openreview.net/forum?id=wqD6TfbYkrn
AUTHORS: Zhaoyang Lyu, Zhifeng Kong, Xudong XU, Liang Pan, Dahua Lin
HIGHLIGHT: To tackle this problem, we propose a novel Point Diffusion-Refinement (PDR) paradigm for point cloud completion.

737, TITLE: Adversarial Robustness Through the Lens of Causality
https://openreview.net/forum?id=cZAi1yWpiXQ
AUTHORS: Yonggang Zhang, Mingming Gong, Tongliang Liu, Gang Niu, Xinmei Tian, Bo Han, Bernhard Sch?lkopf, Kun Zhang
HIGHLIGHT: The first attempt towards using causality to understand and mitigate adversarial vulnerability.

738, TITLE: CDTrans: Cross-domain Transformer for Unsupervised Domain Adaptation
https://openreview.net/forum?id=XGzk5OKWFFc
AUTHORS: Tongkun Xu, Weihua Chen, Pichao WANG, Fan Wang, Hao Li, Rong Jin
HIGHLIGHT: Along with the pseudo labels, a weight-sharing triple-branch transformer framework is proposed to apply self-attention and cross-attention for source/target feature learning and source-target domain alignment, respectively.

739, TITLE: Self-Supervised Graph Neural Networks for Improved Electroencephalographic Seizure Analysis
https://openreview.net/forum?id=k9bx1EfHI_-

AUTHORS: Siyi Tang, Jared Dunnmon, Khaled Kamal Saab, Xuan Zhang, Qianying Huang, Florian Dubost, Daniel Rubin, Christopher Lee-Messer
HIGHLIGHT: Self-supervised graph neural networks for seizure detection and classification from EEG.

740, TITLE: Effective Model Sparsification by Scheduled Grow-and-Prune Methods
https://openreview.net/forum?id=xa6otUDdP2W
AUTHORS: Xiaolong Ma, Minghai Qin, Fei Sun, Zejiang Hou, Kun Yuan, Yi Xu, Yanzhi Wang, Yen-Kuang Chen, Rong Jin, Yuan Xie
HIGHLIGHT: In this paper, we propose a novel scheduled grow-and-prune (GaP) methodology without having to pre-train a dense model.

741, TITLE: Automated Self-Supervised Learning for Graphs
https://openreview.net/forum?id=rFbR4Fv-D6-
AUTHORS: Wei Jin, Xiaorui Liu, Xiangyu Zhao, Yao Ma, Neil Shah, Jiliang Tang
HIGHLIGHT: An automated self-supervised learning algorithm for graph neural networks.

742, TITLE: WeakM3D: Towards Weakly Supervised Monocular 3D Object Detection
https://openreview.net/forum?id=ahi2XSHpAUZ
AUTHORS: Liang Peng, Senbo Yan, Boxi Wu, Zheng Yang, Xiaofei He, Deng Cai
HIGHLIGHT: This paper explores the weakly supervised monocular 3D detection to dispense with the reliance on 3D box labels.

743, TITLE: The Essential Elements of Offline RL via Supervised Learning
https://openreview.net/forum?id=S874XAIpkR-
AUTHORS: Scott Emmons, Benjamin Eysenbach, Ilya Kostrikov, Sergey Levine
HIGHLIGHT: Experimentally evaluating when and why supervised learning solves offline RL

744, TITLE: How Low Can We Go: Trading Memory for Error in Low-Precision Training
https://openreview.net/forum?id=YpSxqy_RE84
AUTHORS: Chengrun Yang, Ziyang Wu, Jerry Chee, Christopher De Sa, Madeleine Udell
HIGHLIGHT: Given a dataset and a memory budget, we use matrix factorization and active learning to efficiently pick the perfect low-precision configuration for a neural network.

745, TITLE: MAML is a Noisy Contrastive Learner
https://openreview.net/forum?id=LDAwu17QaJz
AUTHORS: Chia Hsiang Kao, Wei-Chen Chiu, Pin-Yu Chen
HIGHLIGHT: The Model-agnostic meta learning (MAML) algorithm is a noisy supervised contrastive learner where the noise comes from random initialization and cross-task interference.

746, TITLE: Non-Autoregressive Models are Better Multilingual Translators
https://openreview.net/forum?id=5HvpvYd68b
AUTHORS: Zhenqiao Song, Hao Zhou, Lihua Qian, Jingjing Xu, Shanbo Cheng, Mingxuan Wang, Lei Li
HIGHLIGHT: In this paper, we propose switch-GLAT, a non-autoregressive multilingual machine translation model with a code-switch decoder.

747, TITLE: LEARNING GUARANTEES FOR GRAPH CONVOLUTIONAL NETWORKS ON THE STOCHASTIC BLOCK MODEL
https://openreview.net/forum?id=dpXL6lz4mOQ
AUTHORS: Wei Lu
HIGHLIGHT: In this paper we present the first provable guarantees for one of the best-studied families of graph neural network models, Graph Convolutional Networks (GCNs), for semi- supervised community detection tasks.

748, TITLE: CrossFormer: A Versatile Vision Transformer Hinging on Cross-scale Attention
https://openreview.net/forum?id=_PHymLIxuI
AUTHORS: Wenxiao Wang, Lu Yao, Long Chen, Binbin Lin, Deng Cai, Xiaofei He, Wei Liu
HIGHLIGHT: To this end, we propose Cross-scale Embedding Layer (CEL) and Long Short Distance Attention (LSDA).

749, TITLE: Solving Inverse Problems in Medical Imaging with Score-Based Generative Models
https://openreview.net/forum?id=vaRCHVj0uGI
AUTHORS: Yang Song, Liyue Shen, Lei Xing, Stefano Ermon

HIGHLIGHT: To address this issue, we propose a fully unsupervised technique for inverse problem solving, leveraging the recently introduced score-based generative models.

750, TITLE: Effect of scale on catastrophic forgetting in neural networks
https://openreview.net/forum?id=GhVS8_yPeEa
AUTHORS: Vinay Venkatesh Ramasesh, Aitor Lewkowycz, Ethan Dyer
HIGHLIGHT: We find that large, pre-trained models are robust to catastrophic forgetting.

751, TITLE: CrossBeam: Learning to Search in Bottom-Up Program Synthesis
https://openreview.net/forum?id=qhC8mr2LEKq
AUTHORS: Kensen Shi, Hanjun Dai, Kevin Ellis, Charles Sutton
HIGHLIGHT: We propose training a neural model to learn a hands-on search policy for bottom-up program synthesis, in an effort to tame the search space blowup.

752, TITLE: SQuant: On-the-Fly Data-Free Quantization via Diagonal Hessian Approximation
https://openreview.net/forum?id=JXhROKNZzOc
AUTHORS: Cong Guo, Yuxian Qiu, Jingwen Leng, Xiaotian Gao, Chen Zhang, Yunxin Liu, Fan Yang, Yuhao Zhu, Minyi Guo
HIGHLIGHT: A fast and accurate data-free quantization framework named SQuant.

753, TITLE: W-CTC: a Connectionist Temporal Classification Loss with Wild Cards
https://openreview.net/forum?id=0RqDp8FCW5Z
AUTHORS: Xingyu Cai, Jiahong Yuan, Yuchen Bian, Guangxu Xun, Jiaji Huang, Kenneth Church
HIGHLIGHT: This paper proposes wild-card CTC to solve the problem that the label only matches middle part of the sequence.

754, TITLE: Is Fairness Only Metric Deep? Evaluating and Addressing Subgroup Gaps in Deep Metric Learning
https://openreview.net/forum?id=js62_xuLDDv
AUTHORS: Natalie Dullerud, Karsten Roth, Kimia Hamidieh, Nicolas Papernot, Marzyeh Ghassemi
HIGHLIGHT: We provide a benchmark for fairness in the scope of deep metric learning; investigate fairness impacts of learned representations on downstream classification; and provide a novel method for reducing subgroup gaps in deep metric learning methods.

755, TITLE: Generalized Natural Gradient Flows in Hidden Convex-Concave Games and GANs
https://openreview.net/forum?id=bsycpMi00R1
AUTHORS: Andjela Mladenovic, Iosif Sakos, Gauthier Gidel, Georgios Piliouras
HIGHLIGHT: This paper studies the generalized Gradient Descent-Ascent (GDA) flow in a large class of non-convex non-concave Zero-Sum games dubbed Hidden Convex-Concave games, a class of games that includes GANs.

756, TITLE: ExT5: Towards Extreme Multi-Task Scaling for Transfer Learning
https://openreview.net/forum?id=Vzh1BFUCiIX
AUTHORS: Vamsi Aribandi, Yi Tay, Tal Schuster, Jinfeng Rao, Huaixiu Steven Zheng, Sanket Vaibhav Mehta, Honglei Zhuang, Vinh Q. Tran, Dara Bahri, Jianmo Ni, Jai Gupta, Kai Hui, Sebastian Ruder, Donald Metzler
HIGHLIGHT: Using a suite of 107 NLP tasks, we show that massively multi-task pre-training can improve downstream performance on NLP tasks, overcoming trends of negative transfer between tasks while fine-tuning.

757, TITLE: Multimeasurement Generative Models
https://openreview.net/forum?id=QRX0nCX_gk
AUTHORS: Saeed Saremi, Rupesh Kumar Srivastava
HIGHLIGHT: We formally map the problem of sampling from an unknown distribution with density p_X in $\mathbb{R}^d$ to the problem of learning and sampling $p_\mathbf{Y}$ in $\mathbb{R}^{Md}$ obtained by convolving p_X with a fixed factorial kernel: $p_\mathbf{Y}$ is referred to as M-density and the factorial kernel as multimeasurement noise model (MNM).

758, TITLE: Fair Normalizing Flows
https://openreview.net/forum?id=BrFIKuxrZE
AUTHORS: Mislav Balunovic, Anian Ruoss, Martin Vechev
HIGHLIGHT: We propose a new fair representation learning method based on normalizing flows which can bound the accuracy of any adversary trying to predict sensitive attributes.

759, TITLE: Monotonic Differentiable Sorting Networks
https://openreview.net/forum?id=IcUWShptD7d
AUTHORS: Felix Petersen, Christian Borgelt, Hilde Kuehne, Oliver Deussen
HIGHLIGHT: To address this issue, we propose a novel relaxation of conditional swap operations that guarantees monotonicity in differentiable sorting networks.

760, TITLE: Bridging the Gap: Providing Post-Hoc Symbolic Explanations for Sequential Decision-Making Problems with Inscrutable Representations
https://openreview.net/forum?id=o-1v9hdSult
AUTHORS: Sarath Sreedharan, Utkarsh Soni, Mudit Verma, Siddharth Srivastava, Subbarao Kambhampati
HIGHLIGHT: This paper introduces methods for providing contrastive explanations in terms of user-specified concepts for sequential decision-making settings where the system's model of the task may be best represented as an inscrutable model.

761, TITLE: Diverse Client Selection for Federated Learning via Submodular Maximization
https://openreview.net/forum?id=nwKXyFvaUm
AUTHORS: Ravikumar Balakrishnan, Tian Li, Tianyi Zhou, Nageen Himayat, Virginia Smith, Jeff Bilmes
HIGHLIGHT: The paper addresses a key challenge of selecting the most representative clients iteratively for federated learning through formulating it as a submodular optimization problem and developing efficient algorithms.

762, TITLE: Fast Regression for Structured Inputs
https://openreview.net/forum?id=gNp54NxHUPJ
AUTHORS: Raphael A Meyer, Cameron N Musco, Christopher P Musco, David Woodruff, Samson Zhou
HIGHLIGHT: We show that for a large class of structured inputs, such as combinations of low-rank matrices, sparse matrices, and Vandermonde matrices, regression can be approximately solved using runtime that is polynomial in p.

763, TITLE: Resonance in Weight Space: Covariate Shift Can Drive Divergence of SGD with Momentum
https://openreview.net/forum?id=5ECQL05ub0J
AUTHORS: Kirby Banman, Garnet Liam Peet-Pare, Nidhi Hegde, Alona Fyshe, Martha White
HIGHLIGHT: We show that SGDm under covariate shift with fixed step-size can be unstable and diverge due to a phenomenon known as parametric resonance.

764, TITLE: Audio Lottery: Speech Recognition Made Ultra-Lightweight, Noise-Robust, and Transferable
https://openreview.net/forum?id=9Nk6AJkVYB
AUTHORS: Shaojin Ding, Tianlong Chen, Zhangyang Wang
HIGHLIGHT: We for the first time investigate three unique properties that were rarely studied in previous LTH research but are key to user-interactive ASR devices, bringing new insights to both LTH theory and lightweight ASR research.

765, TITLE: Interacting Contour Stochastic Gradient Langevin Dynamics
https://openreview.net/forum?id=IK9ap6nxXr2
AUTHORS: Wei Deng, Siqi Liang, Botao Hao, Guang Lin, Faming Liang
HIGHLIGHT: We propose an interacting contour stochastic gradient Langevin dynamics sampler and prove it can be theoretically more efficient than a single-chain process with an equivalent computational budget.

766, TITLE: Mapping Language Models to Grounded Conceptual Spaces
https://openreview.net/forum?id=gJcEM8sxHK
AUTHORS: Roma Patel, Ellie Pavlick
HIGHLIGHT: Mapping text-only pre-trained language models to grounded conceptual worlds.

767, TITLE: DeSKO: Stability-Assured Robust Control with a Deep Stochastic Koopman Operator
https://openreview.net/forum?id=hniLRD_XCA
AUTHORS: Minghao Han, Jacob Euler-Rolle, Robert K. Katzschmann
HIGHLIGHT: A robust learning control framework with guarantee stability based on deep stochastic Koopman operator models

768, TITLE: High Probability Generalization Bounds for Minimax Problems with Fast Rates
https://openreview.net/forum?id=gI7feJ9yXPz
AUTHORS: Shaojie Li, Yong Liu
HIGHLIGHT: In this paper, we provide improved generalization analyses for almost all existing generalization measures of minimax problems, which enables the minimax problems to establish sharper bounds of order $\mathcal{O}\left(1/n \right)$, significantly, with high probability.

769, TITLE: Conditional Object-Centric Learning from Video
https://openreview.net/forum?id=aD7uesX1GF_
AUTHORS: Thomas Kipf, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, Klaus Greff
HIGHLIGHT: In this paper, we instead take a weakly-supervised approach and focus on how 1) using the temporal dynamics of video data in the form of optical flow and 2) conditioning the model on simple object location cues can be used to enable segmenting and tracking objects in significantly more realistic synthetic data.

770, TITLE: NETWORK INSENSITIVITY TO PARAMETER NOISE VIA PARAMETER ATTACK DURING TRAINING
https://openreview.net/forum?id=-8sBpe7rDiV
AUTHORS: Julian B?chel, Fynn Firouz Faber, Dylan Richard Muir
HIGHLIGHT: We flatten the weight loss-landscape by introducing a parameter attack term in the loss function and demonstrate improved network insensitivity to noise common in analog neuromorphic hardware.

771, TITLE: Crystal Diffusion Variational Autoencoder for Periodic Material Generation
https://openreview.net/forum?id=03RLpj-tc_
AUTHORS: Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, Tommi S. Jaakkola
HIGHLIGHT: A generative model for the 3D periodic structure of materials
We also provide several standard datasets and evaluation metrics for the broader machine learning community.

772, TITLE: Neural Program Synthesis with Query
https://openreview.net/forum?id=NyJ2KIN8P17
AUTHORS: Di Huang, Rui Zhang, Xing Hu, Xishan Zhang, Pengwei Jin, Nan Li, Zidong Du, Qi Guo, Yunji Chen
HIGHLIGHT: We propose a query-based framework for the interactive program synthesis.

773, TITLE: Learning Scenario Representation for Solving Two-stage Stochastic Integer Programs
https://openreview.net/forum?id=06Wy2BtxXrz
AUTHORS: Yaoxin Wu, Wen Song, Zhiguang Cao, Jie Zhang
HIGHLIGHT: This paper provides a CVAE based method to learn scenario representations for solving stochastic integer programs.

774, TITLE: An Autoregressive Flow Model for 3D Molecular Geometry Generation from Scratch
https://openreview.net/forum?id=C03Ajc-NS5W
AUTHORS: Youzhi Luo, Shuiwang Ji
HIGHLIGHT: We present a novel method for 3D molecular geometry generation from scratch.

775, TITLE: Robbing the Fed: Directly Obtaining Private Data in Federated Learning with Modified Models
https://openreview.net/forum?id=fwzUgo0FM9v
AUTHORS: Liam H Fowl, Jonas Geiping, Wojciech Czaja, Micah Goldblum, Tom Goldstein
HIGHLIGHT: In this work, we introduce a new threat model based on minimal but malicious modifications of the shared model architecture which enable the server to directly obtain a verbatim copy of user data from gradient updates without solving difficult inverse problems.

776, TITLE: Deep Learning without Shortcuts: Shaping the Kernel with Tailored Rectifiers
https://openreview.net/forum?id=U0k7XNTiFEq
AUTHORS: Guodong Zhang, Aleksandar Botev, James Martens
HIGHLIGHT: In this work, we rectify this situation by developing a new type of transformation which is perfectly compatible with a variant of ReLUs -- Leaky ReLUs.

777, TITLE: Learning State Representations via Retracing in Reinforcement Learning
https://openreview.net/forum?id=CLpxpXqqBV
AUTHORS: Changmin Yu, Dong Li, Jianye HAO, Jun Wang, Neil Burgess
HIGHLIGHT: We introduce Learning via Retracing, a novel self-supervised framework based on temporal cycle-consistency assumption of the transition dynamics, for improved learning of the representation (and the dynamics model) in RL tasks.

778, TITLE: Plant 'n' Seek: Can You Find the Winning Ticket?
https://openreview.net/forum?id=9n9c8sf0xm
AUTHORS: Jonas Fischer, Rebekka Burkholz
HIGHLIGHT: We derive a framework to plant ground truth lottery tickets in randomly initialized deep neural networks.

779, TITLE: Dynamic Token Normalization improves Vision Transformers
https://openreview.net/forum?id=f9MHpAGUyMn
AUTHORS: Wenqi Shao, Yixiao Ge, Zhaoyang Zhang, XUYUAN XU, Xiaogang Wang, Ying Shan, Ping Luo
HIGHLIGHT: The proposed DTN is a simple yet effective normalizer for vision transformers.

780, TITLE: When Can We Learn General-Sum Markov Games with a Large Number of Players Sample-Efficiently?
https://openreview.net/forum?id=6MmiS0HUJHR
AUTHORS: Ziang Song, Song Mei, Yu Bai
HIGHLIGHT: We present new algorithms for several learning goals in multi-player general-sum Markov games, with mild PAC sample complexity in terms of the number of players.

781, TITLE: Efficient Computation of Deep Nonlinear Infinite-Width Neural Networks that Learn Features
https://openreview.net/forum?id=tUMr0Iox8XW
AUTHORS: Greg Yang, Michael Santacroce, Edward J Hu
HIGHLIGHT: A new feature learning 8-width limit for deep nonlinear networks closes the performance gap between finite- and infinite-width neural networks previously left by NTK.

782, TITLE: In a Nutshell, the Human Asked for This: Latent Goals for Following Temporal Specifications
https://openreview.net/forum?id=rUwm9wCjURV
AUTHORS: Borja G. Le?n, Murray Shanahan, Francesco Belardinelli
HIGHLIGHT: Inducing architectures to generate low-dimensional representations of their current goal processing observations and instructions together yields stronger out-of-distribution generalisation

783, TITLE: Temporal Alignment Prediction for Supervised Representation Learning and Few-Shot Sequence Classification
https://openreview.net/forum?id=p3DKPQ7uaAi
AUTHORS: Bing Su, Ji-Rong Wen
HIGHLIGHT: We propose a learnable sequence distance by predicting the temporal alignment and show its application in supervised representation learning for sequence data and few-shot action recognition.

784, TITLE: Preference Conditioned Neural Multi-objective Combinatorial Optimization
https://openreview.net/forum?id=QuObT9BTWo
AUTHORS: Xi Lin, Zhiyuan Yang, Qingfu Zhang
HIGHLIGHT: We propose a learning-based method to approximate the whole Pareto set for multi-objective combinatorial optimization problems with a single model.

785, TITLE: Neural Link Prediction with Walk Pooling
https://openreview.net/forum?id=CCu6RcUMwK0
AUTHORS: Liming Pan, Cheng Shi, Ivan Dokmanic
HIGHLIGHT: We propose a link prediction algorithm based on a new pooling scheme called WalkPool.

786, TITLE: The Geometry of Memoryless Stochastic Policy Optimization in Infinite-Horizon POMDPs
https://openreview.net/forum?id=A05I5IvrdL-
AUTHORS: Johannes M?ller, Guido Montufar
HIGHLIGHT: We provide an explicit description of the optimization problem and derive bounds on the number of critical points in POMDPs with memoryless stochastic policies depending on the degree of observability.

787, TITLE: Random matrices in service of ML footprint: ternary random features with no performance loss
https://openreview.net/forum?id=qwULHx9zld
AUTHORS: Hafiz Tiomoko Ali, Zhenyu Liao, Romain Couillet
HIGHLIGHT: A novel computational and storage efficient random features technique with no performance loss

788, TITLE: Generalization of Neural Combinatorial Solvers Through the Lens of Adversarial Robustness
https://openreview.net/forum?id=vJZ7dPIjip3
AUTHORS: Simon Geisler, Johanna Sommer, Jan Schuchardt, Aleksandar Bojchevski, Stephan G?nnemann
HIGHLIGHT: We study the generalization of combinatorial optimization w.r.t. to adversarial attacks since current evaluation protocols are too optimistic and we show that neural solvers are indeed vulnerable under label-preserving perturbations.

789, TITLE: iFlood: A Stable and Effective Regularizer

https://openreview.net/forum?id=MsHnJPaBUZE
AUTHORS: Yuexiang Xie, Zhen WANG, Yaliang Li, Ce Zhang, Jingren Zhou, Bolin Ding
HIGHLIGHT: We propose a novel regularizer named iFlood, which encourages the trained models to better fit the under-fitted instances while suppressing the confidence on over-fitted ones.

790, TITLE: Uncertainty Modeling for Out-of-Distribution Generalization
https://openreview.net/forum?id=6HN7LHyzGgC
AUTHORS: Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, LINGYU DUAN
HIGHLIGHT: We for the first time treat feature statistics as uncertain distributions to improve the model generalization ability.

791, TITLE: ConFeSS: A Framework for Single Source Cross-Domain Few-Shot Learning
https://openreview.net/forum?id=zRJu6mU2BaE
AUTHORS: Debasmit Das, Sungrack Yun, Fatih Porikli
HIGHLIGHT: In this paper, we propose a framework for few-shot learning coined as ConFeSS (Contrastive Learning and Feature Selection System) that tackles large domain shift between base and novel categories.

792, TITLE: Learning with Noisy Labels Revisited: A Study Using Real-World Human Annotations
https://openreview.net/forum?id=TBWA6PLJZQm
AUTHORS: Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, Yang Liu
HIGHLIGHT: In this paper, we revisit the problem of learning from noisy labels using human annotated CIFAR datasets we collected from Amazon Mechanical Turks.

793, TITLE: How Did the Model Change? Efficiently Assessing Machine Learning API Shifts
https://openreview.net/forum?id=gFDFKC4gHL4
AUTHORS: Lingjiao Chen, Matei Zaharia, James Zou
HIGHLIGHT: We systematically study real world ML APIs' performance shifts due to API updates/retraining and propose a framework to efficiently estimate those shifts.

794, TITLE: Auto-Transfer: Learning to Route Transferable Representations
https://openreview.net/forum?id=SIKV0_MrZlr
AUTHORS: Keerthiram Murugesan, Vijay Sadashivaiah, Ronny Luss, Karthikeyan Shanmugam, Pin-Yu Chen, Amit Dhurandhar
HIGHLIGHT: This paper offers a novel transfer method that uses a routing function to select source-target pairs as well as aggregation functions that combine source and target features (as opposed to matching source to target features).

795, TITLE: On the Convergence of Certified Robust Training with Interval Bound Propagation
https://openreview.net/forum?id=YeShU5mLfLt
AUTHORS: Yihan Wang, Zhouxing Shi, Quanquan Gu, Cho-Jui Hsieh
HIGHLIGHT: We present the first theoretical analysis on the convergence of certified robust training with interval bound propagation.

796, TITLE: Minimax Optimality (Probably) Doesn't Imply Distribution Learning for GANs
https://openreview.net/forum?id=nc0ETaieux
AUTHORS: Sitan Chen, Jerry Li, Yuanzhi Li, Raghu Meka
HIGHLIGHT: Under the standard crypto assumption that local pseudorandom generators exist, we show that even a global optimizer for the population WGAN objective need not be close to the true distribution.

797, TITLE: An Information Fusion Approach to Learning with Instance-Dependent Label Noise
https://openreview.net/forum?id=ecH2FKaARUp
AUTHORS: Zhimeng Jiang, Kaixiong Zhou, Zirui Liu, Li Li, Rui Chen, Soo-Hyun Choi, Xia Hu
HIGHLIGHT: This work is the first time to realize and bridge the gap between clean and noisy empirical distribution mismatch.

798, TITLE: Unsupervised Federated Learning is Possible
https://openreview.net/forum?id=WHA8009laxu
AUTHORS: Nan Lu, Zhao Wang, Xiaoxiao Li, Gang Niu, Qi Dou, Masashi Sugiyama
HIGHLIGHT: Federated learning: no label no cry

799, TITLE: Blaschke Product Neural Networks (BPNN): A Physics-Infused Neural Network for Phase Retrieval of Meromorphic Functions
https://openreview.net/forum?id=JJxiD-kg-oK
AUTHORS: Juncheng Dong, Simiao Ren, Yang Deng, Omar Khatib, Jordan Malof, Mohammadreza Soltani, Willie Padilla, Vahid Tarokh
HIGHLIGHT: To this end, we propose a physics-infused deep neural network based on the Blaschke products for phase retrieval.

800, TITLE: Implicit Bias of MSE Gradient Optimization in Underparameterized Neural Networks
https://openreview.net/forum?id=VLgmhQDVBV
AUTHORS: Benjamin Bowman, Guido Montufar
HIGHLIGHT: Underparameterized networks optimizing MSE learn eigenfunctions of an NTK integral operator at rates corresponding to their eigenvalues.

801, TITLE: Semi-relaxed Gromov-Wasserstein divergence and applications on graphs
https://openreview.net/forum?id=RShaMexjc-x
AUTHORS: C?dric Vincent-Cuaz, R?mi Flamary, Marco Corneli, Titouan Vayer, Nicolas Courty
HIGHLIGHT: A new transport based divergence between structured data induced by the relaxation of a mass constraint of the Gromov-Wasserstein problem, leading to new SOTA performances for unsupervised ML applications on graphs.

802, TITLE: Chemical-Reaction-Aware Molecule Representation Learning
https://openreview.net/forum?id=6sh3pIzKS-
AUTHORS: Hongwei Wang, Weijiang Li, Xiaomeng Jin, Kyunghyun Cho, Heng Ji, Jiawei Han, Martin Burke
HIGHLIGHT: We make use of chemical reactions to improve the generalization ability of learned molecule embeddings

803, TITLE: Universalizing Weak Supervision
https://openreview.net/forum?id=YpPiNigTzMT
AUTHORS: Changho Shin, Winfred Li, Harit Vishwakarma, Nicholas Carl Roberts, Frederic Sala
HIGHLIGHT: We extend weak supervision frameworks to new settings ? rankings, regression, Riemannian spaces, and more ? with a universal algorithm with theoretical guarantees.

804, TITLE: Memory Replay with Data Compression for Continual Learning
https://openreview.net/forum?id=a7H7OucbWaU
AUTHORS: Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing HONG, Shifeng Zhang, Zhenguo Li, Yi Zhong, Jun Zhu
HIGHLIGHT: We propose memory replay with data compression, which is an important yet neglected baseline and a promising direction for continual learning.

805, TITLE: Cold Brew: Distilling Graph Node Representations with Incomplete or Missing Neighborhoods
https://openreview.net/forum?id=1ugNpm7W6E
AUTHORS: Wenqing Zheng, Edward W Huang, Nikhil Rao, Sumeet Katariya, Zhangyang Wang, Karthik Subbian
HIGHLIGHT: Improve strict cold start performances for graph minings with a knowledge distillation framework.

806, TITLE: Graph-Guided Network for Irregularly Sampled Multivariate Time Series
https://openreview.net/forum?id=Kwm8I7dU-l5
AUTHORS: Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, Marinka Zitnik
HIGHLIGHT: Here, we introduce RAINDROP, a graph-guided network for learning representations of irregularly sampled multivariate time series.

807, TITLE: Transformer-based Transform Coding
https://openreview.net/forum?id=IDwN6xjHnK8
AUTHORS: Yinhao Zhu, Yang Yang, Taco Cohen
HIGHLIGHT: Specifically, we show that nonlinear transforms built on Swin-transformers can achieve better compression efficiency than transforms built on convolutional neural networks (ConvNets), while requiring fewer parameters and shorter decoding time.

808, TITLE: On the benefits of maximum likelihood estimation for Regression and Forecasting
https://openreview.net/forum?id=zrW-LVXj2k1
AUTHORS: Pranjal Awasthi, Abhimanyu Das, Rajat Sen, Ananda Theertha Suresh
HIGHLIGHT: MLE + Post-Hoc Inference can be competitive with any estimator under some general assumptions

809, TITLE: VC dimension of partially quantized neural networks in the overparametrized regime
https://openreview.net/forum?id=7udZAsEzd60
AUTHORS: Yutong Wang, Clayton Scott
HIGHLIGHT: We apply VC theory to analyze the performance of a neural network in the overparametrized regime and obtain a minimax-optimality result.

810, TITLE: ARTEMIS: Attention-based Retrieval with Text-Explicit Matching and Implicit Similarity
https://openreview.net/forum?id=CVfLvQq9gLo
AUTHORS: Ginger Delmas, Rafael S. Rezende, Gabriela Csurka, Diane Larlus
HIGHLIGHT: Our work aims at shedding new light on the task by looking at it through the prism of two familiar and related frameworks: text-to-image and image-to-image retrieval.

811, TITLE: Which Shortcut Cues Will DNNs Choose? A Study from the Parameter-Space Perspective
https://openreview.net/forum?id=qRDQi3ocgR3
AUTHORS: Luca Scimeca, Seong Joon Oh, Sanghyuk Chun, Michael Poli, Sangdoo Yun
HIGHLIGHT: When given equally likely shortcuts in data, which shortcut cue will a DNN choose, and why?

812, TITLE: RelViT: Concept-guided Vision Transformer for Visual Relational Reasoning
https://openreview.net/forum?id=afoV8W3-IYp
AUTHORS: Xiaojian Ma, Weili Nie, Zhiding Yu, Huaizu Jiang, Chaowei Xiao, Yuke Zhu, Song-Chun Zhu, Anima Anandkumar
HIGHLIGHT: We propose a novel concept-feature dictionary to enable two new concept-guided auxiliary tasks, which largely improve the model performances on visual relational reasoning, especially for systematic generalization.

813, TITLE: Demystifying Batch Normalization in ReLU Networks: Equivalent Convex Optimization Models and Implicit Regularization
https://openreview.net/forum?id=6XGgutacQ0B
AUTHORS: Tolga Ergen, Arda Sahiner, Batu Ozturkler, John M. Pauly, Morteza Mardani, Mert Pilanci
HIGHLIGHT: We introduce an analytic framework based on convex duality to obtain exact and polynomial-time trainable convex representations of weight-decay regularized ReLU networks with BN.

814, TITLE: Towards Continual Knowledge Learning of Language Models
https://openreview.net/forum?id=vfsRB5MImo9
AUTHORS: Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun KIM, Stanley Jungkyu Choi, Minjoon Seo
HIGHLIGHT: We propose a novel continual learning formulation named Continual Knowledge Learning which allows large language models to constantly obtain new and updated knowledge while mitigating forgetting of previous learned time-invariant knowledge.
We construct a new benchmark and metric to quantify the retention of time-invariant world knowledge, the update of outdated knowledge, and the acquisition of new knowledge.

815, TITLE: On Covariate Shift of Latent Confounders in Imitation and Reinforcement Learning
https://openreview.net/forum?id=w01vBAcewNX
AUTHORS: Guy Tennenholtz, Assaf Hallak, Gal Dalal, Shie Mannor, Gal Chechik, Uri Shalit
HIGHLIGHT: We use expert data with unobserved confounders for both imitation and reinforcement learning. Such hidden confounding is prone to a shifted distribution, which may severely hurt performance unless accounted for.

816, TITLE: Graph Auto-Encoder via Neighborhood Wasserstein Reconstruction
https://openreview.net/forum?id=ATUh28lnSuW
AUTHORS: Mingyue Tang, Pan Li, Carl Yang
HIGHLIGHT: We study unsupervised graph representation learning and propose a novel decoder based on neighborhood reconstruction with Wasserstein distance to facilitate the GNN encoding of entire neighborhood information beyond direct links.

817, TITLE: Complete Verification via Multi-Neuron Relaxation Guided Branch-and-Bound
https://openreview.net/forum?id=l_amHf1oaK
AUTHORS: Claudio Ferrari, Mark Niklas Mueller, Nikola Jovanovic, Martin Vechev
HIGHLIGHT: We obtain a state-of-the-art GPU-based neural network verifier by leveraging tight multi-neuron constraints in a Branch-and-Bound setting.

818, TITLE: Non-Linear Operator Approximations for Initial Value Problems

https://openreview.net/forum?id=d2TT6gK9qZn
AUTHORS: Gaurav Gupta, Xiongye Xiao, Radu Balan, Paul Bogdan
HIGHLIGHT: A Pad? approximation based exponential operator module is proposed for working with the Initial Value Problems. The compactness of model yields data-efficiency and better performance which is demonstrated on scarce real-world dataset.

819, TITLE: Query Embedding on Hyper-Relational Knowledge Graphs
https://openreview.net/forum?id=4rLw09TgRw9
AUTHORS: Dimitrios Alivanistos, Max Berrendorf, Michael Cochez, Mikhail Galkin
HIGHLIGHT: We investigate how to extend the multi-hop reasoning problem to hyper-relational queries on knowledge graphs and consider methods for solving it.

820, TITLE: IGLU: Efficient GCN Training via Lazy Updates
https://openreview.net/forum?id=5kq11Tl1z4
AUTHORS: S Deepak Narayanan, Aditya Sinha, Prateek Jain, Purushottam Kar, SUNDARARAJAN SELLAMANICKAM
HIGHLIGHT: IGLU is a novel lazy update-based optimization technique for accelerated GCN training with provable convergence guarantees

821, TITLE: DKM: Differentiable k-Means Clustering Layer for Neural Network Compression
https://openreview.net/forum?id=J_F_qqCE3Z5
AUTHORS: Minsik Cho, Keivan Alizadeh-Vahid, Saurabh Adya, Mohammad Rastegari
HIGHLIGHT: We propose a novel model compression scheme based on differentiable K-means layer, and it delivers the state-of-the-art results.

822, TITLE: A Generalized Weighted Optimization Method for Computational Learning and Inversion
https://openreview.net/forum?id=14F3fI6MGxX
AUTHORS: Kui Ren, Yunan Yang, Bj?rn Engquist
HIGHLIGHT: This paper proposes a generalized weighted optimization method for computational learning and inversion with noisy data and derives generalization error bounds for various feature regression models, demonstrating better generalization capabilities.

823, TITLE: A Neural Tangent Kernel Perspective of Infinite Tree Ensembles
https://openreview.net/forum?id=vUH85MOXO7h
AUTHORS: Ryuichi Kanoh, Mahito Sugiyama
HIGHLIGHT: By considering an ensemble of infinite trees, we introduce and study the Tree Neural Tangent Kernel (TNTK), which provides new insights into the behavior of the infinite ensemble of soft trees.

824, TITLE: Learning Graphon Mean Field Games and Approximate Nash Equilibria
https://openreview.net/forum?id=0sgntlpKDOz
AUTHORS: Kai Cui, Heinz Koeppl
HIGHLIGHT: We propose, analyze and solve a novel, theoretically well-founded graph-based mean field game for Nash equilibria in discrete-time dynamical systems on otherwise infeasibly large dense graphs.

825, TITLE: Differentiable DAG Sampling
https://openreview.net/forum?id=9wOQOgNe-w
AUTHORS: Bertrand Charpentier, Simon Kibler, Stephan G?nnemann
HIGHLIGHT: We propose a new differentiable probabilistic model over DAGs (DP-DAG).

826, TITLE: Online Coreset Selection for Rehearsal-based Continual Learning
https://openreview.net/forum?id=f9D-5WNG4Nv
AUTHORS: Jaehong Yoon, Divyam Madaan, Eunho Yang, Sung Ju Hwang
HIGHLIGHT: We propose Online Coreset Selection (OCS), a simple yet effective method that selects the most representative and informative coreset at each iteration and trains them in an online manner.

827, TITLE: Large-Scale Representation Learning on Graphs via Bootstrapping
https://openreview.net/forum?id=0UXT6PpRpW
AUTHORS: Shantanu Thakoor, Corentin Tallec, Mohammad Gheshlaghi Azar, Mehdi Azabou, Eva L Dyer, Remi Munos, Petar Velickovic, Michal Valko
HIGHLIGHT: To address these challenges, we introduce Bootstrapped Graph Latents (BGRL) - a graph representation learning method that learns by predicting alternative augmentations of the input.

828, TITLE: Trust Region Policy Optimisation in Multi-Agent Reinforcement Learning
https://openreview.net/forum?id=EcGGFkNTxdJ
AUTHORS: Jakub Grudzien Kuba, Ruiqing Chen, Muning Wen, Ying Wen, Fanglei Sun, Jun Wang, Yaodong Yang
HIGHLIGHT: This paper introduces the first trust region method for multi-agent reinforcement learning that enjoys theoretically-justified monotonic improvement guarantee and demonstrates the state-of-the-art performance on Mujoco benchmarks.

829, TITLE: A generalization of the randomized singular value decomposition
https://openreview.net/forum?id=hgKtwSb4S2
AUTHORS: Nicolas Boulle, Alex Townsend
HIGHLIGHT: The randomized SVD is generalized to multivariate Gaussian input vectors and Hilbert-Schmidt operators.

830, TITLE: Understanding and Leveraging Overparameterization in Recursive Value Estimation
https://openreview.net/forum?id=shbAgEsk3qM
AUTHORS: Chenjun Xiao, Bo Dai, Jincheng Mei, Oscar A Ramirez, Ramki Gummadi, Chris Harris, Dale Schuurmans
HIGHLIGHT: We present an analysis of value estimation under overparameterized linear representations, and develop new algorithmic tools for improving recursive value estimation with deep models based on the new findings.

831, TITLE: Open-World Semi-Supervised Learning
https://openreview.net/forum?id=O-r8LOR-CCA
AUTHORS: Kaidi Cao, Maria Brbic, Jure Leskovec
HIGHLIGHT: We propose a pipeline that recognizes previously seen classes and discovers novel, never-before-seen classes at the same time.

832, TITLE: AlphaZero-based Proof Cost Network to Aid Game Solving
https://openreview.net/forum?id=nKWjE4QF1hB
AUTHORS: Ti-Rong Wu, Chung-Chin Shih, Ting Han Wei, Meng-Yu Tsai, Wei-Yuan Hsu, I-Chen Wu
HIGHLIGHT: This paper proposes a novel approach to solving problems by modifying the training target of the AlphaZero algorithm, such that it prioritizes solving the game quickly, rather than winning.

833, TITLE: A Reduction-Based Framework for Conservative Bandits and Reinforcement Learning
https://openreview.net/forum?id=AcrlgZ9BKed
AUTHORS: Yunchang Yang, Tianhao Wu, Han Zhong, Evrard Garcelon, Matteo Pirotta, Alessandro Lazaric, Liwei Wang, Simon Shaolei Du
HIGHLIGHT: We give general framework that turns upper and lower bounds in non-conservative settings to bounds in conservative settings.

834, TITLE: Improving Mutual Information Estimation with Annealed and Energy-Based Bounds
https://openreview.net/forum?id=T0B9AoM_bFg
AUTHORS: Rob Brekelmans, Sicong Huang, Marzyeh Ghassemi, Greg Ver Steeg, Roger Baker Grosse, Alireza Makhzani
HIGHLIGHT: We derive new annealed importance sampling and energy-based bounds, resulting in vastly more accurate estimates of mutual information.

835, TITLE: The Convex Geometry of Backpropagation: Neural Network Gradient Flows Converge to Extreme Points of the Dual Convex Program
https://openreview.net/forum?id=5QhUE1qiVC6
AUTHORS: Yifei Wang, Mert Pilanci
HIGHLIGHT: We study non-convex subgradient flows for training two-layer ReLU neural networks from a convex geometry and duality perspective.

836, TITLE: NAS-Bench-Suite: NAS Evaluation is (Now) Surprisingly Easy
https://openreview.net/forum?id=0DLwqQLmqV
AUTHORS: Yash Mehta, Colin White, Arber Zela, Arjun Krishnakumar, Guri Zabergja, Shakiba Moradian, Mahmoud Safari, Kaicheng Yu, Frank Hutter
HIGHLIGHT: We show that you cannot get away only with NAS-Bench-101 and -201; to fix this, we release a unified NAS benchmark suite with 25 benchmarks.

837, TITLE: Anomaly Detection for Tabular Data with Internal Contrastive Learning
https://openreview.net/forum?id=_hszZbt46bT
AUTHORS: Tom Shenkar, Lior Wolf
HIGHLIGHT: An anomaly detection method based on the ability to predict the masked out part in a vector.

838, TITLE: The Boltzmann Policy Distribution: Accounting for Systematic Suboptimality in Human Models
https://openreview.net/forum?id=_1_QjPGN5ye
AUTHORS: Cassidy Laidlaw, Anca Dragan
HIGHLIGHT: We propose modeling human behavior with a Boltzmann distribution over policies?not trajectories?and show it is more accurate and useful.

839, TITLE: Quantitative Performance Assessment of CNN Units via Topological Entropy Calculation
https://openreview.net/forum?id=xFOyMwWPkz
AUTHORS: Yang Zhao, Hao Zhang
HIGHLIGHT: We propose a novel method for quantitatively clarifying the status of individual units in CNNs and show its value in interpreting networks with different generalization ability.

840, TITLE: Who Is the Strongest Enemy? Towards Optimal and Efficient Evasion Attacks in Deep RL
https://openreview.net/forum?id=JM2kFbJvvI
AUTHORS: Yanchao Sun, Ruijie Zheng, Yongyuan Liang, Furong Huang
HIGHLIGHT: We theoretically characterize the essence of evasion attacks in RL, and propose a novel attack algorithm for RL agents, which achieves state-of-the-art performance on both attacking and robustifying RL agents in many Atari and MuJoCo tasks.

841, TITLE: FastSHAP: Real-Time Shapley Value Estimation
https://openreview.net/forum?id=Zq2G_VTV53T
AUTHORS: Neil Jethani, Mukund Sudarshan, Ian Connick Covert, Su-In Lee, Rajesh Ranganath
HIGHLIGHT: We introduce FastSHAP, a new method for estimating Shapley values in a single forward pass using an explainer model that is learned via stochastic gradient optimization using a weighted least squares-like objective function.

842, TITLE: FairCal: Fairness Calibration for Face Verification
https://openreview.net/forum?id=nRj0NcmSuxb
AUTHORS: Tiago Salvador, Stephanie Cairns, Vikram Voleti, Noah Marshall, Adam M Oberman
HIGHLIGHT: We calibrate face verification models for fairness, without use of the sensitive attribute, without the need for retraining. This leads to SOTA accuracy, fairness calibration, and equal FPRs across subgroups.

843, TITLE: A Loss Curvature Perspective on Training Instabilities of Deep Learning Models
https://openreview.net/forum?id=OcKMT-36vUs
AUTHORS: Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Edward Dahl, Zachary Nado, Orhan Firat
HIGHLIGHT: Our results suggest a unifying perspective on how disparate mitigation strategies for training instability ultimately address poor conditioning.

844, TITLE: GradMax: Growing Neural Networks using Gradient Information
https://openreview.net/forum?id=qjN4h_wwUO
AUTHORS: Utku Evci, Bart van Merrienboer, Thomas Unterthiner, Fabian Pedregosa, Max Vladymyrov
HIGHLIGHT: We present a method that adds new neurons during training without impacting what is already learned, while improving the training dynamics.

845, TITLE: Learning Object-Oriented Dynamics for Planning from Text
https://openreview.net/forum?id=B6EIcyp-Rb7
AUTHORS: Guiliang Liu, Ashutosh Adhikari, Amir-massoud Farahmand, Pascal Poupart
HIGHLIGHT: We propose an Object-Oriented Text Dynamic (OOTD) model for solving decision-making problems in the text domain.

846, TITLE: On-Policy Model Errors in Reinforcement Learning
https://openreview.net/forum?id=81e1aeOt-sd
AUTHORS: Lukas Froehlich, Maksym Lefarov, Melanie Zeilinger, Felix Berkenkamp
HIGHLIGHT: We combine real-world data and a learned model for data-efficient reinforcement learning with reduced model-bias.

847, TITLE: Learning to Dequantise with Truncated Flows
https://openreview.net/forum?id=fExcSKdDo_
AUTHORS: Shawn Tan, Chin-Wei Huang, Alessandro Sordoni, Aaron Courville
HIGHLIGHT: Learning a variational dequantisation scheme with truncated/bounded-support distributions

848, TITLE: Offline Neural Contextual Bandits: Pessimism, Optimization and Generalization
https://openreview.net/forum?id=sPIFuucA3F
AUTHORS: Thanh Nguyen-Tang, Sunil Gupta, A. Tuan Nguyen, Svetha Venkatesh
HIGHLIGHT: Provably efficient offline contextual bandits with neural network function approximation

849, TITLE: Learning Features with Parameter-Free Layers
https://openreview.net/forum?id=bCrdi4iVvv
AUTHORS: Dongyoon Han, YoungJoon Yoo, Beomyoung Kim, Byeongho Heo
HIGHLIGHT: This paper introduces a new design paradigm rethinking a parameter-free operation as the main building block of network architecture.

850, TITLE: Who Is Your Right Mixup Partner in Positive and Unlabeled Learning
https://openreview.net/forum?id=NH29920YEmj
AUTHORS: Changchun Li, Ximing Li, Lei Feng, Jihong Ouyang
HIGHLIGHT: We propose a novel PU learning method named P3Mix which simultaneously benefits from instance augmentation and supervision correction with a heuristic mixup technique.

851, TITLE: Generalizing Few-Shot NAS with Gradient Matching
https://openreview.net/forum?id=_jMtny3sMKU
AUTHORS: Shoukang Hu, Ruochen Wang, Lanqing HONG, Zhenguo Li, Cho-Jui Hsieh, Jiashi Feng
HIGHLIGHT: In this work, we propose a gradient matching score (GM) that leverages gradient information at the shared weight for making informed splitting decisions.

852, TITLE: Generating Videos with Dynamics-aware Implicit Generative Adversarial Networks
https://openreview.net/forum?id=Czsdv-S4-w9
AUTHORS: Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, Jinwoo Shin
HIGHLIGHT: We make video generation scalable leveraging implicit neural representations.

853, TITLE: CoordX: Accelerating Implicit Neural Representation with a Split MLP Architecture
https://openreview.net/forum?id=oAy7yPmdNz
AUTHORS: Ruofan Liang, Hongyi Sun, Nandita Vijaykumar
HIGHLIGHT: In this work, we aim to accelerate inference and training of coordinate-based MLPs for implicit neural representations by proposing a new split MLP architecture, CoordX.

854, TITLE: Graph-Relational Domain Adaptation
https://openreview.net/forum?id=kcwyXtt7yDJ
AUTHORS: Zihao Xu, Hao He, Guang-He Lee, Bernie Wang, Hao Wang
HIGHLIGHT: In this work, we relax such uniform alignment by using a domain graph to encode domain adjacency, e.g., a graph of states in the US with each state as a domain and each edge indicating adjacency, thereby allowing domains to align flexibly based on the graph structure.

855, TITLE: MetaShift: A Dataset of Datasets for Evaluating Contextual Distribution Shifts and Training Conflicts
https://openreview.net/forum?id=MTex8qKavoS
AUTHORS: Weixin Liang, James Zou
HIGHLIGHT: We leverage annotated subsets within a heterogeneous dataset to evaluate the performance of learning algorithms to distribution shifts and to visualize training dynamics.

856, TITLE: Zero-CL: Instance and Feature decorrelation for negative-free symmetric contrastive learning
https://openreview.net/forum?id=RAW9tCdVxLj
AUTHORS: Shaofeng Zhang, Feng Zhu, Junchi Yan, Rui Zhao, Xiaokang Yang
HIGHLIGHT: We develop two contrastive learning methods to prevent collapses in symmetric architecture without negative pairs.

857, TITLE: Understanding Intrinsic Robustness Using Label Uncertainty
https://openreview.net/forum?id=6ET9SzlgNX
AUTHORS: Xiao Zhang, David Evans
HIGHLIGHT: We argue that the standard concentration fails to fully characterize the intrinsic robustness of a classification problem, since it ignores data labels which are essential to any classification task.

858, TITLE: Evaluating Distributional Distortion in Neural Language Modeling
https://openreview.net/forum?id=bTteFbU99ye
AUTHORS: Benjamin LeBrun, Alessandro Sordoni, Timothy J. O'Donnell
HIGHLIGHT: To address this gap, we develop a controlled evaluation scheme which uses generative models trained on natural data as artificial languages from which we can exactly compute sequence probabilities.

859, TITLE: Learning to Schedule Learning rate with Graph Neural Networks
https://openreview.net/forum?id=k7efTb0un9z
AUTHORS: Yuanhao Xiong, Li-Cheng Lan, Xiangning Chen, Ruochen Wang, Cho-Jui Hsieh
HIGHLIGHT: We propose a novel Graph-Network-based Scheduler (GNS), which is both informative to to encode rich information and generalizable to different architectures.

860, TITLE: Omni-Scale CNNs: a simple and effective kernel size configuration for time series classification
https://openreview.net/forum?id=PDYs7Z2XFGv
AUTHORS: Wensi Tang, Guodong Long, Lu Liu, Tianyi Zhou, Michael Blumenstein, Jing Jiang
HIGHLIGHT: To extract features from time series data in proper time scales, many complicated scales searching or weighting methods have been proposed, but we will show that this could have been done via a very simple structure.

861, TITLE: Post-Training Detection of Backdoor Attacks for Two-Class and Multi-Attack Scenarios
https://openreview.net/forum?id=MSgB8D4Hy51
AUTHORS: Zhen Xiang, David Miller, George Kesidis
HIGHLIGHT: We proposed a detection framework against backdoor attacks for two-class and multi-attack scenarios, without access to the classifier's training set or any supervision from clean classifiers trained for the same domain.

862, TITLE: MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer
https://openreview.net/forum?id=vh-0sUt8HlG
AUTHORS: Sachin Mehta, Mohammad Rastegari
HIGHLIGHT: Light-weight and general-purpose vision transformers for mobile devices

863, TITLE: Unraveling Model-Agnostic Meta-Learning via The Adaptation Learning Rate
https://openreview.net/forum?id=3rULBvOJ8D2
AUTHORS: Yingtian Zou, Fusheng Liu, Qianxiao Li
HIGHLIGHT: Theoretical analysis of Model-Agnostic Meta-Learning (MAML) through the inner loop (adaptation) learning rate.

864, TITLE: DAB-DETR: Dynamic Anchor Boxes are Better Queries for DETR
https://openreview.net/forum?id=oMI9PjOb9Jl
AUTHORS: Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, Lei Zhang
HIGHLIGHT: We present in this paper a novel query formulation using dynamic anchor boxes for DETR and offer a deeper understanding of the role of queries in DETR.

865, TITLE: Subspace Regularizers for Few-Shot Class Incremental Learning
https://openreview.net/forum?id=boJy41J-tnQ
AUTHORS: Afra Feyza Aky?rek, Ekin Aky?rek, Derry Wijaya, Jacob Andreas
HIGHLIGHT: We propose a simple yet highly effective set of subspace-based regularizers to address representation learning for few-shot incremental classification.

866, TITLE: Learn Locally, Correct Globally: A Distributed Algorithm for Training Graph Neural Networks
https://openreview.net/forum?id=FndDxSz3LxQ
AUTHORS: Morteza Ramezani, Weilin Cong, Mehrdad Mahdavi, Mahmut Kandemir, Anand Sivasubramaniam
HIGHLIGHT: We propose LLCG a communication efficient distributed algorithm for training GNNs.

867, TITLE: Unsupervised Semantic Segmentation by Distilling Feature Correspondences
https://openreview.net/forum?id=SaKO6z6Hl0c
AUTHORS: Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, William T. Freeman
HIGHLIGHT: We use the correlations between self-supervised visual features to perform unsupervised semantic segmentation.

868, TITLE: Hidden Convexity of Wasserstein GANs: Interpretable Generative Models with Closed-Form Solutions

https://openreview.net/forum?id=e2Lle5cij9D
AUTHORS: Arda Sahiner, Tolga Ergen, Batu Ozturkler, Burak Bartan, John M. Pauly, Morteza Mardani, Mert Pilanci
HIGHLIGHT: We demonstrate that Wasserstein GANs with two-layer discriminators and a variety of generators are equivalent to convex optimization problems or convex-concave games, allowing for global optimization in polynomial time and improved interpretability.

869, TITLE: Autonomous Reinforcement Learning: Formalism and Benchmarking
https://openreview.net/forum?id=nkaba3ND7B5
AUTHORS: Archit Sharma, Kelvin Xu, Nikhil Sardana, Abhishek Gupta, Karol Hausman, Sergey Levine, Chelsea Finn
HIGHLIGHT: In this paper, we aim to address this discrepancy by laying out a framework for Autonomous Reinforcement Learning (ARL): reinforcement learning where the agent not only learns through its own experience, but also contends with lack of human supervision to reset between trials.

870, TITLE: CADDA: Class-wise Automatic Differentiable Data Augmentation for EEG Signals
https://openreview.net/forum?id=6IYp-35L-xJ
AUTHORS: C?dric Rommel, Thomas Moreau, Joseph Paillard, Alexandre Gramfort
HIGHLIGHT: We propose a new gradient-based automatic data augmentation technique where samples are transformed based on their labels, and demonstrate its relevance on EEG sleep staging data.

871, TITLE: Declarative nets that are equilibrium models
https://openreview.net/forum?id=q4HaTeMO--y
AUTHORS: Russell Tsuchida, Suk Yee Yong, Mohammad Ali Armin, Lars Petersson, Cheng Soon Ong
HIGHLIGHT: Choosing a kernelised generalised linear model as the inner problem of a DDN yields a DEQ with specific fixed weights.

872, TITLE: The Role of Pretrained Representations for the OOD Generalization of RL Agents
https://openreview.net/forum?id=8eb12UQYxrG
AUTHORS: Frederik Tr?uble, Andrea Dittadi, Manuel Wuthrich, Felix Widmaier, Peter Vincent Gehler, Ole Winther, Francesco Locatello, Olivier Bachem, Bernhard Sch?lkopf, Stefan Bauer
HIGHLIGHT: We study the role of pretrained representations for the out-of-distribution generalization of RL agents.

873, TITLE: On Incorporating Inductive Biases into VAEs
https://openreview.net/forum?id=nzvbBD_3J-g
AUTHORS: Ning Miao, Emile Mathieu, Siddharth N, Yee Whye Teh, Tom Rainforth
HIGHLIGHT: A flexible and effective framework for adding inductive biases to VAEs that corrects the pathologies of previous approaches and leads to improved representations and generative models.

874, TITLE: Simple GNN Regularisation for 3D Molecular Property Prediction and Beyond
https://openreview.net/forum?id=1wVvweK3oIb
AUTHORS: Jonathan Godwin, Michael Schaarschmidt, Alexander L Gaunt, Alvaro Sanchez-Gonzalez, Yulia Rubanova, Petar Velickovic, James Kirkpatrick, Peter Battaglia
HIGHLIGHT: A simple regularisation technique for GNNs applied to 3D molecular property prediction & beyond.

875, TITLE: Auto-scaling Vision Transformers without Training
https://openreview.net/forum?id=H94a1_Pyr-6
AUTHORS: Wuyang Chen, Wei Huang, Xianzhi Du, Xiaodan Song, Zhangyang Wang, Denny Zhou
HIGHLIGHT: We automate the design and scaling of vision transformers without any training, achieving state-of-the-art performance on ImageNet classification and COCO object detection.

876, TITLE: Robust Learning Meets Generative Models: Can Proxy Distributions Improve Adversarial Robustness?
https://openreview.net/forum?id=WVX0NNVBBkV
AUTHORS: Vikash Sehwag, Saeed Mahloujifar, Tinashe Handina, Sihui Dai, Chong Xiang, Mung Chiang, Prateek Mittal
HIGHLIGHT: We leverage proxy distributions to significantly improve the robustness of deep neural network.

877, TITLE: Generative Modeling with Optimal Transport Maps
https://openreview.net/forum?id=5JdLZg346Lw
AUTHORS: Litu Rout, Alexander Korotin, Evgeny Burnaev
HIGHLIGHT: While optimal transport cost serves as the loss for popular generative models, we demonstrate that the optimal transport map can be used as the generative model itself.

878, TITLE: EigenGame Unloaded: When playing games is better than optimizing
https://openreview.net/forum?id=So6YAqnqgMj
AUTHORS: Ian Gemp, Brian McWilliams, Claire Vernade, Thore Graepel
HIGHLIGHT: We improve the EigenGame algorithm by removing update bias, enabling further parallelism and better performance.

879, TITLE: Convergent and Efficient Deep Q Learning Algorithm
https://openreview.net/forum?id=OJm3HZuj4r7
AUTHORS: Zhikang T. Wang, Masahito Ueda
HIGHLIGHT: In this work, we show that DQN can indeed diverge and cease to operate in realistic settings.

880, TITLE: Differentially Private Fine-tuning of Language Models
https://openreview.net/forum?id=Q42f0dfjECO
AUTHORS: Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, Sergey Yekhanin, Huishuai Zhang
HIGHLIGHT: We show that by combining recent advances in NLP, parameter-efficiency, privacy accounting, and using larger models, one can privately fine-tune models whose utility approaches that of non-private models.

881, TITLE: Spread Spurious Attribute: Improving Worst-group Accuracy with Spurious Attribute Estimation
https://openreview.net/forum?id=_F9xpOrqyX9
AUTHORS: Junhyun Nam, Jaehyung Kim, Jaeho Lee, Jinwoo Shin
HIGHLIGHT: Using a small amount of attribute annotated samples for training can boost worst-group performance in the presence of spurious correlation.

882, TITLE: Exploring extreme parameter compression for pre-trained language models
https://openreview.net/forum?id=RftryyYyjiG
AUTHORS: Benyou Wang, Yuxin Ren, Lifeng Shang, Xin Jiang, Qun Liu
HIGHLIGHT: In this work, we aim to explore larger compression ratios for PLMs, among which tensor decomposition is a potential but under-investigated one.

883, TITLE: Exploring Memorization in Adversarial Training
https://openreview.net/forum?id=7gE9V9GBZaI
AUTHORS: Yinpeng Dong, Ke Xu, Xiao Yang, Tianyu Pang, Zhijie Deng, Hang Su, Jun Zhu
HIGHLIGHT: This paper explores the memorization effect in adversarial training and analyzes its connections with model capacity, convergence, generalization, and especially robust overfitting of the adversarially trained models.

884, TITLE: Decoupled Adaptation for Cross-Domain Object Detection
https://openreview.net/forum?id=VNqaB1g9393
AUTHORS: Junguang Jiang, Baixu Chen, Jianmin Wang, Mingsheng Long
HIGHLIGHT: To deal with the challenges in cross-domain object detection, we propose D-adapt to decouple the adversarial adaptation and the training of the detector, and also decouple the category adaptation and the bounding box adaptation.

885, TITLE: Sequential Reptile: Inter-Task Gradient Alignment for Multilingual Learning
https://openreview.net/forum?id=ivQruZvXxtz
AUTHORS: Seanie Lee, Hae Beom Lee, Juho Lee, Sung Ju Hwang
HIGHLIGHT: We propose a simple yet effective gradient alignment method for finetuning multilingual pretrained language models.

886, TITLE: GRAND++: Graph Neural Diffusion with A Source Term
https://openreview.net/forum?id=EMxu-dzvJk
AUTHORS: Matthew Thorpe, Tan Minh Nguyen, Hedi Xia, Thomas Strohmer, Andrea Bertozzi, Stanley Osher, Bao Wang
HIGHLIGHT: We propose GRAND++ for deep graph learning with limited labeled training data

887, TITLE: High Probability Bounds for a Class of Nonconvex Algorithms with AdaGrad Stepsize
https://openreview.net/forum?id=dSw0QtRMJkO
AUTHORS: Ali Kavis, Kfir Yehuda Levy, Volkan Cevher
HIGHLIGHT: In this paper, we propose a new, simplified high probability analysis of AdaGrad for smooth, non-convex problems.

888, TITLE: FedPara: Low-rank Hadamard Product for Communication-Efficient Federated Learning

https://openreview.net/forum?id=d71n4ftoCBy
AUTHORS: Nam Hyeon-Woo, Moon Ye-Bin, Tae-Hyun Oh
HIGHLIGHT: New communication-efficient neural network parameterization for federated learning.

889, TITLE: Online Facility Location with Predictions
https://openreview.net/forum?id=DSQHjibtgKR
AUTHORS: Shaofeng H.-C. Jiang, Erzhi Liu, You Lyu, Zhihao Gavin Tang, Yubo Zhang
HIGHLIGHT: We give a nearly optimal robust algorithm for online facility location with predictions.

890, TITLE: Training Transition Policies via Distribution Matching for Complex Tasks
https://openreview.net/forum?id=6vkzF28Hur8
AUTHORS: JU-SEUNG BYUN, Andrew Perrault
HIGHLIGHT: Training transition policies via distribution matching

891, TITLE: Neural Markov Controlled SDE: Stochastic Optimization for Continuous-Time Data
https://openreview.net/forum?id=7DI6op61AY
AUTHORS: Sung Woo Park, Kyungjae Lee, Junseok Kwon
HIGHLIGHT: We propose a novel probabilistic framework for modelling stochastic dynamics with the rigorous use of optimal control theory.

892, TITLE: Mirror Descent Policy Optimization
https://openreview.net/forum?id=aBO5SvgSt1
AUTHORS: Manan Tomar, Lior Shani, Yonathan Efroni, Mohammad Ghavamzadeh
HIGHLIGHT: A theory-grounded practical algorithm for policy optimization in RL, which is conceptually simpler and performs better or on par to SOTA.

893, TITLE: Sample Selection with Uncertainty of Losses for Learning with Noisy Labels
https://openreview.net/forum?id=xENf4QUL4LW
AUTHORS: Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Jun Yu, Gang Niu, Masashi Sugiyama
HIGHLIGHT: In this paper, we incorporate the uncertainty of losses by adopting interval estimation instead of point estimation of losses, where lower bounds of the con?dence intervals of losses derived from distribution-free concentration inequalities, but not losses themselves, are used for sample selection.

894, TITLE: Causal ImageNet: How to discover spurious features in Deep Learning?
https://openreview.net/forum?id=XVPqLyNxSyh
AUTHORS: Sahil Singla, Soheil Feizi
HIGHLIGHT: A scalable framework for discovering spurious features of deep neural networks
Using this methodology, we introduce the {\it Causal Imagenet} dataset containing causal and spurious masks for a large set of samples from Imagenet.

895, TITLE: Towards Distribution Shift of Node-Level Prediction on Graphs: An Invariance Perspective
https://openreview.net/forum?id=FQOC5u-1egI
AUTHORS: Qitian Wu, Hengrui Zhang, Junchi Yan, David Wipf
HIGHLIGHT: We formulate out-of-distribution generalization problem for node-level prediction on graphs and propose a new learning approach based on invariant models

896, TITLE: A Deep Variational Approach to Clustering Survival Data
https://openreview.net/forum?id=RQ428ZptQfU
AUTHORS: Laura Manduchi, Ricards Marcinkevics, Michela C. Massi, Thomas Weikert, Alexander Sauter, Verena Gotta, Timothy M?ller, Flavio Vasella, Marian C. Neidert, Marc Pfister, Bram Stieltjes, Julia E Vogt
HIGHLIGHT: We introduce a novel semi-supervised probabilistic approach to cluster survival data

897, TITLE: VAT-Mart: Learning Visual Action Trajectory Proposals for Manipulating 3D ARTiculated Objects
https://openreview.net/forum?id=iEx3PiooLy
AUTHORS: Ruihai Wu, Yan Zhao, Kaichun Mo, Zizheng Guo, Yian Wang, Tianhao Wu, Qingnan Fan, Xuelin Chen, Leonidas Guibas, Hao Dong
HIGHLIGHT: We propose a novel interaction-for-perception framework to learn visual actionable representations (i.e. affordance and action trajectory proposals) for robotic manipulation.

898, TITLE: Scale Mixtures of Neural Network Gaussian Processes

https://openreview.net/forum?id=YVPBh4k78iZ
AUTHORS: Hyungi Lee, Eunggu Yun, Hongseok Yang, Juho Lee
HIGHLIGHT: Infinitely-wide neural networks can be equivalent to scale mixtures of Gaussian processes.

899, TITLE: The Three Stages of Learning Dynamics in High-dimensional Kernel Methods
https://openreview.net/forum?id=EQmAP4F859
AUTHORS: Nikhil Ghosh, Song Mei, Bin Yu
HIGHLIGHT: We study the training dynamics of gradient flows on the population risk and empirical risk of high-dimensional kernel least-squares problems, which we show has three learning stages.

900, TITLE: On the Learning of Quasimetrics
https://openreview.net/forum?id=y0VvIg25yk
AUTHORS: Tongzhou Wang, Phillip Isola
HIGHLIGHT: We theoretically analyze various algorithms on learning quasimetrics (asymmetrical metrics), and propose an embedding-based method with strong guarantees. Experiments on graph learning and Q-learning show its effectiveness over common baselines.

901, TITLE: Transferable Visual Control Policies Through Robot-Awareness
https://openreview.net/forum?id=o0ehFykKVtr
AUTHORS: Edward S. Hu, Kun Huang, Oleh Rybkin, Dinesh Jayaraman
HIGHLIGHT: We closely integrate readily available knowledge about the robot and world into a learned model to facilitate transfer.

902, TITLE: Focus on the Common Good: Group Distributional Robustness Follows
https://openreview.net/forum?id=irARV_2VFs4
AUTHORS: Vihari Piratla, Praneeth Netrapalli, Sunita Sarawagi
HIGHLIGHT: We propose a new and simple algorithm for the sub-population shift problem that enables learning of shared features and performed consistently well over several standard, and real-world, benchmarks of the problem.

903, TITLE: Minimax Optimization with Smooth Algorithmic Adversaries
https://openreview.net/forum?id=UdxJ2fJx7N0
AUTHORS: Tanner Fiez, Chi Jin, Praneeth Netrapalli, Lillian J Ratliff
HIGHLIGHT: We propose a tractable formulation of minimax optimization by modeling the adversary's algorithm, and present new algorithms which are guaranteed to converge and find appropriate stationary points.

904, TITLE: Node Feature Extraction by Self-Supervised Multi-scale Neighborhood Prediction
https://openreview.net/forum?id=KJggliHbs8
AUTHORS: Eli Chien, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Jiong Zhang, Olgica Milenkovic, Inderjit S Dhillon
HIGHLIGHT: We design a self-supervised learning method for extracting node representations from raw data.

905, TITLE: New Insights on Reducing Abrupt Representation Change in Online Continual Learning
https://openreview.net/forum?id=N8MaByOzUfb
AUTHORS: Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, Eugene Belilovsky
HIGHLIGHT: We study how representations shift at task boundaries in the single-head online continual learning setting, leading to a simple high performance method

906, TITLE: Proving the Lottery Ticket Hypothesis for Convolutional Neural Networks
https://openreview.net/forum?id=Vjki79-619-
AUTHORS: Arthur da Cunha, Emanuele Natale, Laurent Viennot
HIGHLIGHT: We prove the lottery ticket hypothesis for convolutional neural networks

907, TITLE: Generative Principal Component Analysis
https://openreview.net/forum?id=pgir5f7ekAL
AUTHORS: Zhaoqiang Liu, Jiulong Liu, Subhroshekhar Ghosh, Jun Han, Jonathan Scarlett
HIGHLIGHT: We study the problem of principal component analysis with generative modeling assumptions, and provide a corresponding efficient algorithm with provable guarantees.

908, TITLE: Mastering Visual Continuous Control: Improved Data-Augmented Reinforcement Learning
https://openreview.net/forum?id=_SJ-_yyes8

AUTHORS: Denis Yarats, Rob Fergus, Alessandro Lazaric, Lerrel Pinto
HIGHLIGHT: We proposed a model-free off-policy algorithm for image-based continuous control that significantly outperforms previous methods both in sample and time complexity.

909, TITLE: Fast AdvProp
https://openreview.net/forum?id=hcoswsDHNAW
AUTHORS: Jieru Mei, Yucheng Han, Yutong Bai, Yixiao Zhang, Yingwei Li, Xianhang Li, Alan Yuille, Cihang Xie
HIGHLIGHT: In this paper, we introduce Fast AdvProp, which aggressively revamps AdvProp's costly training components, rendering the method nearly as cheap as the vanilla training setting.

910, TITLE: Fixed Neural Network Steganography: Train the images, not the network
https://openreview.net/forum?id=hcMvApxGSzZ
AUTHORS: Varsha Kishore, Xiangyu Chen, Yan Wang, Boyi Li, Kilian Q Weinberger
HIGHLIGHT: A novel method for steganography based on adversarial perturbations.

911, TITLE: Distribution Compression in Near-Linear Time
https://openreview.net/forum?id=lzupY5zjaU9
AUTHORS: Abhishek Shetty, Raaz Dwivedi, Lester Mackey
HIGHLIGHT: We introduce a simple algorithm for compressing an n-point summary of a probability distribution into a $\sqrt{n}$-point summary of comparable quality in $O(n \log^2 n)$ time.

912, TITLE: AdaMatch: A Unified Approach to Semi-Supervised Learning and Domain Adaptation
https://openreview.net/forum?id=Q5uh1Nvv5dm
AUTHORS: David Berthelot, Rebecca Roelofs, Kihyuk Sohn, Nicholas Carlini, Alexey Kurakin
HIGHLIGHT: We introduce AdaMatch, a unified solution that achieves state-of-the-art results for unsupervised domain adaptation (UDA), semi-supervised learning (SSL), and semi-supervised domain adaptation (SSDA).

913, TITLE: Feature Kernel Distillation
https://openreview.net/forum?id=tBIQEvApZK5
AUTHORS: Bobby He, Mete Ozay
HIGHLIGHT: A feature-learning perspective of (ensemble) Knowledge Distillation (KD) in Neural Networks to propose a new method (FKD), with both theoretical & experimental results demonstrating FKD's advantages over standard KD baselines.

914, TITLE: Back2Future: Leveraging Backfill Dynamics for Improving Real-time Predictions in Future
https://openreview.net/forum?id=L01Nn_VJ9i
AUTHORS: Harshavardhan Kamarthi, Alexander Rodr?guez, B. Aditya Prakash
HIGHLIGHT: We study the problem of multi-variate backfill for both features and targets and show how to leverage our insights for more general neural framework to improve both model predictions and evaluation
We construct a detailed dataset composed of relevant signals over the past year of the pandemic.

915, TITLE: How to Train Your MAML to Excel in Few-Shot Classification
https://openreview.net/forum?id=49h_IkpJtaE
AUTHORS: Han-Jia Ye, Wei-Lun Chao
HIGHLIGHT: In this paper, we point out several key facets of how to train MAML to excel in few-shot classification.

916, TITLE: Dual Lottery Ticket Hypothesis
https://openreview.net/forum?id=fOsN52jn25l
AUTHORS: Yue Bai, Huan Wang, ZHIQIANG TAO, Kunpeng Li, Yun Fu
HIGHLIGHT: We articulate a Dual Lottery Ticket Hypothesis (DLTH) with a proposed training strategy Random Sparse Network to validate DLTH.

917, TITLE: Path Integral Sampler: A Stochastic Control Approach For Sampling
https://openreview.net/forum?id=_uCb2ynRu7Y
AUTHORS: Qinsheng Zhang, Yongxin Chen
HIGHLIGHT: We present Path Integral Sampler~(PIS), an efficient algorithm to draw samples from unnormalized probability density functions.

918, TITLE: DIVA: Dataset Derivative of a Learning Task
https://openreview.net/forum?id=bVvMOtLMiw
AUTHORS: Yonatan Dukler, Alessandro Achille, Giovanni Paolini, Avinash Ravichandran, Marzia Polito, Stefano Soatto

HIGHLIGHT: Presents a method to optimize a dataset based on a notion of a dataset derivative that is computed in closed form using linearization

919, TITLE: Hindsight is 20/20: Leveraging Past Traversals to Aid 3D Perception
https://openreview.net/forum?id=qsZoGvFiJn1
AUTHORS: Yurong You, Katie Z Luo, Xiangyu Chen, Junan Chen, Wei-Lun Chao, Wen Sun, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger
HIGHLIGHT: To this end, we propose a novel end-to-end trainable Hindsight framework to extract this contextual information from past traversals and store it in an easy-to-query data structure, which can then be leveraged to aid future 3D object detection of the same scene.

920, TITLE: Improving Non-Autoregressive Translation Models Without Distillation
https://openreview.net/forum?id=I2Hw58KHp8O
AUTHORS: Xiao Shi Huang, Felipe Perez, Maksims Volkovs
HIGHLIGHT: Improving the CMLM non-autoregressive machine translation model so it trains without knowledge distillation and achieves SOTA BLEU score on both raw and distilled dataset

921, TITLE: FP-DETR: Detection Transformer Advanced by Fully Pre-training
https://openreview.net/forum?id=yjMQuLLcGWK
AUTHORS: Wen Wang, Yang Cao, Jing Zhang, Dacheng Tao
HIGHLIGHT: To mitigate the issue, we propose FP-DETR, a new method that Fully Pre-Trains an encoder-only transformer and smoothly fine-tunes it for object detection via a task adapter.

922, TITLE: Pix2seq: A Language Modeling Framework for Object Detection
https://openreview.net/forum?id=e42KbIw6Wb
AUTHORS: Ting Chen, Saurabh Saxena, Lala Li, David J. Fleet, Geoffrey Hinton
HIGHLIGHT: We demonstrated that object detection can be tackled by simply training a language model conditioned on pixel inputs.

923, TITLE: Is Importance Weighting Incompatible with Interpolating Classifiers?
https://openreview.net/forum?id=uqBOne3LUKy
AUTHORS: Ke Alexander Wang, Niladri Shekhar Chatterji, Saminul Haque, Tatsunori Hashimoto
HIGHLIGHT: We theoretically and empirically demonstrate that importance weighting can be effective in handling distribution shifts in overparameterized classifiers.

924, TITLE: MoReL: Multi-omics Relational Learning
https://openreview.net/forum?id=DnG75_KyHjX
AUTHORS: Arman Hasanzadeh, Ehsan Hajiramezanali, Nick Duffield, Xiaoning Qian
HIGHLIGHT: We propose a novel deep Bayesian generative model to efficiently infer a multi-partite graph encoding molecular interactions across such heterogeneous views, using a fused Gromov-Wasserstein (FGW) regularization between latent representations of corresponding views for integrative analysis.

925, TITLE: Learning to Map for Active Semantic Goal Navigation
https://openreview.net/forum?id=swrMQttr6wN
AUTHORS: Georgios Georgakis, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Kostas Daniilidis
HIGHLIGHT: A framework for object goal navigation that actively learns to predict semantic maps and choose long-term goals based on uncertainty measures.

926, TITLE: Autonomous Learning of Object-Centric Abstractions for High-Level Planning
https://openreview.net/forum?id=rrWeE9ZDw_
AUTHORS: Steven James, Benjamin Rosman, George Konidaris
HIGHLIGHT: We learn object-centric PDDL representations directly from raw observation data

927, TITLE: Concurrent Adversarial Learning for Large-Batch Training
https://openreview.net/forum?id=rw1mZl_ss3L
AUTHORS: Yong Liu, Xiangning Chen, Minhao Cheng, Cho-Jui Hsieh, Yang You
HIGHLIGHT: In this paper, we propose to use adversarial learning to increase the batch size in large-batch training.

928, TITLE: Fine-grained Differentiable Physics: A Yarn-level Model for Fabrics
https://openreview.net/forum?id=KPEFXR1HdIo

AUTHORS: Deshan Gong, Zhanxing Zhu, Andrew J. Bulpitt, He Wang
HIGHLIGHT: Following this motivation, we propose a new differentiable fabrics model for composite materials such as cloths, where we dive into the granularity of yarns and model individual yarn physics and yarn-to-yarn interactions.

929, TITLE: Enabling Arbitrary Translation Objectives with Adaptive Tree Search
https://openreview.net/forum?id=rhOiUS8KQM9
AUTHORS: Wang Ling, Wojciech Stokowiec, Domenic Donato, Chris Dyer, Lei Yu, Laurent Sartran, Austin Matthews
HIGHLIGHT: MCTS is used as a decoder for autoregressive and non-autoregressive machine translation models.

930, TITLE: ClimateGAN: Raising Climate Change Awareness by Generating Images of Floods
https://openreview.net/forum?id=EZNOb_uNpJk
AUTHORS: Victor Schmidt, Alexandra Luccioni, M?lisande Teng, Tianyu Zhang, Alexia Reynaud, Sunand Raghupathi, Gautier Cosne, Adrien Juraver, Vahe Vardanyan, Alex Hern?ndez-Garc?a, Yoshua Bengio
HIGHLIGHT: This paper presents a model to robustly produce photo-realistic images of floods for raising climate change awareness, leveraging unsupervised domain adaptation and conditional image generation.

931, TITLE: LoRA: Low-Rank Adaptation of Large Language Models
https://openreview.net/forum?id=nZeVKeeFYf9
AUTHORS: Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen
HIGHLIGHT: Finetuning updates have a low "intrinsic rank" which allows us to train only the rank decomposition matrices of certain weights, yielding better performance and practical benefits.

932, TITLE: Online Target Q-learning with Reverse Experience Replay: Efficiently finding the Optimal Policy for Linear MDPs
https://openreview.net/forum?id=HMJdXzbWKH
AUTHORS: Naman Agarwal, Syomantak Chaudhuri, Prateek Jain, Dheeraj Mysore Nagaraj, Praneeth Netrapalli
HIGHLIGHT: The goal of this work is to bridge the gap between practical success of Q-learning and the relatively pessimistic theoretical results.

933, TITLE: Memory Augmented Optimizers for Deep Learning
https://openreview.net/forum?id=NRX9QZ6yqt
AUTHORS: Paul-Aymeric Martin McRae, Prasanna Parthasarathi, Mido Assran, Sarath Chandar
HIGHLIGHT: We propose a framework of memory augmented optimizers and empirically show that the class of optimizers provide accelerated convergence and even better test performance. We show the proposed optimizers converge in smooth strongly convex setting.

934, TITLE: Permutation Compressors for Provably Faster Distributed Nonconvex Optimization
https://openreview.net/forum?id=GugZ5DzzAu
AUTHORS: Rafal Szlendak, Alexander Tyurin, Peter Richt?rik
HIGHLIGHT: In this paper, we present the novel compression scheme for distributed non-convex optimization.

935, TITLE: Scene Transformer: A unified architecture for predicting future trajectories of multiple agents
https://openreview.net/forum?id=Wm3EA5OlHsG
AUTHORS: Jiquan Ngiam, Vijay Vasudevan, Benjamin Caine, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, David J Weiss, Ben Sapp, Zhifeng Chen, Jonathon Shlens
HIGHLIGHT: We introduce a scene-centric masked sequence based motion prediction model that unifies a variety of motion prediction tasks from joint motion predictions to conditioned prediction.

936, TITLE: Learning to Extend Molecular Scaffolds with Structural Motifs
https://openreview.net/forum?id=ZTsoE8G3GG
AUTHORS: Krzysztof Maziarz, Henry Richard Jackson-Flux, Pashmina Cameron, Finton Sirockin, Nadine Schneider, Nikolaus Stiefl, Marwin Segler, Marc Brockschmidt
HIGHLIGHT: We propose a new fragment-based generative model of molecules that can be constrained to include an arbitrary subgraph (scaffold).

937, TITLE: Neural Relational Inference with Node-Specific Information
https://openreview.net/forum?id=HBsJNesj2S
AUTHORS: Ershad Banijamali
HIGHLIGHT: We use variational inference to uncover relations among agents in a multi-agent system, given that the agents can have access to some private information

938, TITLE: Demystifying Limited Adversarial Transferability in Automatic Speech Recognition Systems
https://openreview.net/forum?id=l5aSHXi8jG5
AUTHORS: Hadi Abdullah, Aditya Karlekar, Vincent Bindschaedler, Patrick Traynor
HIGHLIGHT: Uncover factors that limit transferability of the popular optimization attacks in the automatic speech recognition systems.

939, TITLE: Pessimistic Model-based Offline Reinforcement Learning under Partial Coverage
https://openreview.net/forum?id=tyrJsbKAe6
AUTHORS: Masatoshi Uehara, Wen Sun
HIGHLIGHT: We study model-based offline Reinforcement Learning with general function approximation without a full coverage assumption on the offline data distribution.

940, TITLE: Learning Towards The Largest Margins
https://openreview.net/forum?id=hqkhcFHOeKD
AUTHORS: Xiong Zhou, Xianming Liu, Deming Zhai, Junjun Jiang, Xin Gao, Xiangyang Ji
HIGHLIGHT: In this work, we attempt to address this limitation by formulating the principled optimization objective as learning towards the largest margins.

941, TITLE: Inverse Online Learning: Understanding Non-Stationary and Reactionary Policies
https://openreview.net/forum?id=DYypjaRdph2
AUTHORS: Alex Chan, Alicia Curth, Mihaela van der Schaar
HIGHLIGHT: We introduce a practical algorithm for retrospectively estimating such perceived effects, alongside the process through which agents update them, using a novel architecture built upon an expressive family of deep state-space models.

942, TITLE: Rethinking Class-Prior Estimation for Positive-Unlabeled Learning
https://openreview.net/forum?id=aYAA-XHKyk
AUTHORS: Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Gang Niu, Masashi Sugiyama, Dacheng Tao
HIGHLIGHT: Class-Prior Estimation for Positive-Unlabeled Learning

943, TITLE: On the Importance of Difficulty Calibration in Membership Inference Attacks
https://openreview.net/forum?id=3eIrli0TwQ
AUTHORS: Lauren Watson, Chuan Guo, Graham Cormode, Alexandre Sablayrolles
HIGHLIGHT: Membership inference attacks can greatly benefit from a technique called difficulty calibration, significantly improving their reliability.

944, TITLE: You are AllSet: A Multiset Function Framework for Hypergraph Neural Networks
https://openreview.net/forum?id=hpBTIv2uy_E
AUTHORS: Eli Chien, Chao Pan, Jianhao Peng, Olgica Milenkovic
HIGHLIGHT: We propose a multiset function framework for hypergraph neural networks.

945, TITLE: Train Short, Test Long: Attention with Linear Biases Enables Input Length Extrapolation
https://openreview.net/forum?id=R8sQPpGCv0
AUTHORS: Ofir Press, Noah Smith, Mike Lewis
HIGHLIGHT: We show that our simple position method enables transformer LMs to efficiently and accurately perform inference on longer sequences than they were trained on.