



36th Annual Conference of Neural Information Processing Systems (NeurIPS)

First Hybrid Program

2022 Fact Sheet

Location: United States, Louisiana, New Orleans Ernest N. Morial Convention Center

Total Hybrid Registrations: 15,530

- Number of in-person registrations: 9,560
- Number of virtual registrations: 5,970

Previous Locations:

- 17,091 Virtual Conference 2021
- 22,823 First Virtual Conference 2020
- 13,000 Vancouver, British Columbia, Canada 2019
- 8,648 Montreal, Quebec, Canada 2018
- 8,008 Long Beach, California, United States 2017
- 5,231 Barcelona, Spain 2016
- 3,852 Montreal, Quebec, Canada 2015
- 2,581 Montreal, Quebec, Canada 2014
- 1,994 Lake Tahoe, California, United States 2013
- 1,676 Lake Tahoe, California, United States 2012
- 1,452 Granada, Spain 2011
- 1,354 Vancouver, British Columbia, Canada 2010

2022 Hybrid Program organized to maximize interaction and engagement:

- 28 November - 3 December (Week 1) in-person program:
 - 7 Invited Talks "Keynotes"
 - 62 Workshops
 - 10 Affinity Group sessions
 - 16 Socials
 - 6 two-hour Poster sessions
 - 49 posters Journal Track Papers
- 5-9 December (Week 2) virtual program
 - 13 [Tutorials](#)
 - Workshops continue
 - 20 posters Journal Track Papers
- 76 Panels
- 26 Competitions

2022 Papers Broken Down:

- 2,905 accepted [papers](#)
- 9,634 full paper submissions
- 20% paper acceptance percentage

- 13 Prestigious Paper Award Winners
- 2 Datasets and Benchmark Award Winners
- Test of Time Award Winner
- 184 Main Orals
- 16 Datasets and Benchmark Orals
- 25 papers in the [Competition Track](#)
- 10,406 Reviewers
- 1,000 Top Reviewers
- Highlighted paper topics
 - neural networks
 - reinforce learning
 - language models
 - graph neural
 - federated learning
 - representation learning
 - general model
 - deep learning
 - vision transformation
 - offline reinforcement

Program Features:

- (New) [Educational Outreach](#) program
 - 247 High School Students, 16 Teachers from 12 Local Schools
 - 87 Volunteers
- (New) [Journal Track Papers](#) program
 - 69 papers presented
 - 49 in person
 - 20 virtually
 - 31 ReScience journal papers
 - 38 JMLR journal papers
- Second annual [Dataset and Benchmarks Track](#)
 - 447 submitted papers for dataset and benchmark track compared to 484 in 2021
 - 163 accepted papers for dataset and benchmark track compared to 174 in 2021
 - 36.46% acceptance rate
 - 92 Area Chairs
 - 1064 Reviewers, 39 Ethic reviewers
- Expo
 - 77 Exhibitors
 - [Expo Talk Panels](#), [Expo Workshops](#) and [Expo Demonstrations](#)

2021 Program stats for comparison:

- 8 conference tracks
- 2,334 papers accepted
- 9,122 paper submissions
- 11,292 abstract submissions
- 25.6 % paper acceptance
- 60 Workshops
- 10 Tutorials
- 22 Competitions
- 18 Demonstrations
- 55 Orals
- 15 Socials
- 40 virtual and local Meetups

2022 Organizing Committee:

- 742 Area Chairs
- 82 Super Area Chairs
- **Program chairs:**
Alekh Agarwal, Google Research
Danielle Belgrave, Deepmind
Kyunghyun Cho, Genentech, NYU
Alice Oh, KAIST
- **General chairs:**
Sanmi Koyejo, Stanford, Google Research
Shakir Mohamed, Deepmind

Affinity Groups Represented:

- Black in AI
- Global South in AI
- Indigenous in AI
- LatinX in AI
- North Africans in ML
- Queer in AI
- Women in ML (WiML)

Invited Keynote Speakers:

- Rediet Abebe - [Algorithms On the Bench: Examining Validity of ML Systems in the Public Sphere](#), Tuesday 29 November 9:30 am CST
- Emmanuel Candes - [Conformal Prediction in 2022](#) - Tuesday, 29 November 2:30 pm CST
- David Chalmers - [Are Large Language Models Sentient?](#) - Monday, 28 November 5:15 PM CST
- Isabelle Guyon - [The Data-Centric Era: How ML is Becoming an Experimental Science](#) - Thursday, 1 December 9:30 am CST
- Geoffrey Hinton, 2022 Test of Time Award Winner - [The Forward-Forward Algorithm for Training Deep Neural Networks](#) - Thursday, 1 December, 2:30 pm CST (*Virtual - pwd: fishvale*)
- Juho Kim - [Interaction-Centric AI](#) - Wednesday, 30 November, 9:30 am CST
- Alondra Nelson - [Blueprint for an AI Bill of Rights Making Automated Systems Work for the American People](#) - Wednesday, 30 November, 2:30 pm CST

Thirteen Outstanding Paper Award Recipients:

[Is Out-of-distribution Detection Learnable?](#)

by Zhen Fang, Yixuan Li, Jie Lu, Jiahua Dong, Bo Han, Feng Liu

This work provides a theoretical study of out-of-distribution (OOD) detection, focusing on the conditions under which such models are learnable. The work uses probably approximately correct (PAC) learning theory to show that OOD detection models are PAC learnable only for some conditions of the space of data distributions and the space of prediction models. It provides 3 concrete impossibility theorems, which can be easily applied to determine the feasibility of OOD detection in practical settings, and which was used in this work to provide a theoretical grounding for existing OOD detection approaches. This work also raises new theoretical questions, for example, about the learnability of near-OOD detection. As such, it has the potential for broad theoretical and practical impact in this important research area.

Tuesday, November 29 — [Poster Session 1](#)

[Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding](#)

by Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Raphael Gontijo-Lopes, Tim Salimans, Jonathan Ho, David J Fleet, Mohammad Norouzi

High quality generative models of images based on Diffusion Process are having a huge impact both within and beyond machine learning. This work represents one of the state of the art of such models, but also innovates in demonstrating the effective combination of an independently trained large language model with an image decoder at scale. This inherently practical decoupling is likely to be a dominant paradigm for large scale text to image models. The results are impressive and of interest to a broad audience.

Thursday, December 1 — [Poster Session 5](#)

[Elucidating the Design Space of Diffusion-Based Generative Models](#)

by Tero Karras, Miika Aittala, Timo Aila, Samuli Laine

This paper is an excellent demonstration of how a well thought through survey, that seeks not just to list but to organise prior research into a coherent common framework, can provide insights that then lead to new modeling improvements. In this case the focus of this paper are generative models of images that incorporate some form of Diffusion Process, which have become extremely popular recently despite the difficulties of training such models. This paper is likely to be an important contribution in the evolution of both the understanding and implementation of Diffusion Process based models.

Wednesday, December 7 — [Featured Papers Panels 3B](#)

[ProcTHOR: Large-Scale Embodied AI Using Procedural Generation](#)

by Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, Roozbeh Mottaghi

This work provides a framework for training embodied AI agents on large quantities of data, creating the potential for such agents to benefit from scaling, as language and image generation models have. The core of the framework is an engine for building procedurally-generated, physics-enabled environments with which agents can interact. This engine, in combination with provided digital assets and environmental controls, allows for generating a combinatorially large number of diverse environments. The authors demonstrate that this framework can be used to train SoTA models for several embodied AI

tasks. The framework and code used in this work will be open-sourced, providing a valuable asset for the research community.

Wednesday, November 30 — Poster Session 3

[Using natural language and program abstractions to instill human inductive biases in machines](#)

by Sreejan Kumar, Carlos G Correa, Ishita Dasgupta, Raja Marjeh, Michael Hu, Robert D. Hawkins, Jonathan Cohen, Nathaniel Daw, Karthik R Narasimhan, Thomas L. Griffiths

Co-training on program abstractions and natural language enables incorporating human biases into learning. This is a clean approach to incorporating human biases but also be robust with program abstractions.

Thursday, December 1 — Poster Session 6

[A Neural Corpus Indexer for Document Retrieval](#)

by Yujing Wang, Yingyan Hou, Haonan Wang, Ziming Miao, Shibin Wu, Hao Sun, Qi Chen, Yuqing Xia, Chengmin Chi, Guoshuai Zhao, Zheng Liu, Xing Xie, Hao Sun, Weiwei Deng, Qi Zhang, Mao Yang

This work proposes a neural indexer that takes as input a query and outputs, via a decoder combined with beam search, a list of IDs corresponding to relevant documents in the index. It joins a small but growing line of research that departs from the dominant high recall-sparse retrieval paradigm. Notably, this new paradigm allows for gradient-based optimization of the indexer for target applications using standard deep learning algorithms and frameworks. The proposed approach introduces architectural and training choices that result in significant improvements compared to prior work, demonstrating the promise of neural indexers as a viable alternative. The paper is well-written and discusses the limitations and open questions following from this work, which can serve as inspiration for future research.

Thursday, December 1 — Poster Session 5

[High-dimensional limit theorems for SGD: Effective dynamics and critical scaling](#)

by Gerard Ben Arous, Reza Gheissari, Aukosh Jagannath

This work studies the scaling limits of SGD with constant step-size in the high-dimensional regime. It shows how complex SGD can be if the step size is large. Characterizing the nature of SDE and comparing it to the ODE when the step size is small gives insights into the nonconvex optimization landscape.

[Gradient Descent: The Ultimate Optimizer](#)

by Kartik Chandra, Audrey Xie, Jonathan Ragan-Kelley, Erik Meijer

This paper reduces sensitivity to hyperparameters in gradient descent by developing a method to optimize with respect to hyperparameters and recursively optimize *hyper*-hyperparameters. Since gradient descent is everywhere, the potential impact is tremendous.

Wednesday, November 30 — Poster Session 4

[Riemannian Score-Based Generative Modelling](#)

by Valentin De Bortoli, Emile Mathieu, Michael John Hutchinson, James Thornton, Yee Whye Teh, Arnaud Doucet

The paper generalizes score-based generative model (SGM) from Euclidean space to Riemannian manifolds by identifying major components that contribute to the success of SGMs. The method is both

a novel and technically useful contribution.

Wednesday, November 30 — Poster Session 4

[Gradient Estimation with Discrete Stein Operators](#)

by *Jiaxin Shi, Yuhao Zhou, Jessica Hwang, Michalis Titsias, Lester Mackey*

This paper considers gradient estimation when the distribution is discrete. Most common gradient estimators suffer from excessive variance. To improve the quality of gradient estimation, they introduce a variance reduction technique based on Stein operators for discrete distributions. Even though Stein operator is classical, this work provides a nice interpretation of it for gradient estimation and also shows practical improvement in experiments.

Tuesday, November 29 — Poster Session 1

[An empirical analysis of compute-optimal large language model training](#)

by *Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack William Rae, Laurent Sifre*

The work asks “Given a fixed FLOPs budget, how should one trade-off model size and the number of training tokens?”. The work models this trade off, makes a prediction based on this model, and trains a model corresponding to that prediction. The resultant model, that is significantly smaller but is trained on significantly more tokens, outperforms its counterpart, while also being more practical to use downstream due to its smaller size. All in all, this work sheds new light on the way the community thinks about scale in the context of language models, which may be useful in other domains of AI as well.

Wednesday, November 30 — Poster Session 4

[Beyond neural scaling laws: beating power law scaling via data pruning](#)

by *Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, Ari S. Morcos*

The importance of high quality data in order to achieve good results in machine learning is well known. Recent work on scaling laws has treated data quality as uniform and focussed on the relationship between computation and data. This work renews our focus on the importance of selecting high quality data as a means to achieve optimal scaling. It does so through a nicely designed analytic investigation that develops a theoretical model of the impact of data quality in concert with empirical instantiation of a range of data filtering metrics on ImageNet. This work is both insightful and timely and will shape the debate about the tradeoffs in the many dimensions of scale in machine learning.

Wednesday, November 30 — Poster Session 3

[On-Demand Sampling: Learning Optimally from Multiple Distributions](#)

by *Nika Haghtalab, Michael Jordan, Eric Zhao*

This paper studies multiple distribution learning using techniques from stochastic zero-sum games. This technique leads to very interesting theoretical results for a class of problems with near optimal results.

Wednesday, November 30 — Poster Session 3

Two Outstanding Datasets and Benchmarks Papers:

[LAION-5B: An open large-scale dataset for training next generation image-text models](#)

by Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, Jenia Jitsev

Studying the training and capabilities of language-vision architectures, such as CLIP and DALL-E, requires datasets containing billions of image-text pairs. Until now, no datasets of this size have been made openly available for the broader research community. This work presents LAION-5B, a dataset consisting of 5.85 billion CLIP-filtered image-text pairs, aimed at democratizing research on large-scale multi-modal models. Moreover, the authors use this data to successfully replicate foundational models such as CLIP, GLIDE and Stable Diffusion, provide several nearest neighbor indices, as well as an improved web-interface, and detection scores for watermark, NSFW, and toxic content detection. Wednesday, November 30 — Poster Session 4

[MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge](#)

by Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, Anima Anandkumar

Autonomous agents have made great strides in specialist domains like Atari games and Go, but typically fail to generalize across a wide spectrum of tasks and capabilities. This work introduces MineDojo, a new framework built on the popular Minecraft game that features a simulation suite with thousands of diverse open-ended tasks and an internet-scale knowledge base with Minecraft videos, tutorials, wiki pages, and forum discussions. It also proposes a novel agent learning algorithm that is able to solve a variety of open-ended tasks specified in free-form language. It provides an open-source simulation suite, knowledge bases, algorithm implementation, and pretrained models to promote research on generally capable embodied agents. Tuesday, November 29 — Poster Session 2

Test of Time Award:

A NeurIPS paper was selected from 10 years ago, “[ImageNet Classification with Deep Convolutional Neural Networks](#)” by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, aka “AlexNet paper” which was unanimously selected by the Program Chairs. In 2012, it was presented as the first CNN trained on the ImageNet Challenge, far surpassing the state-of-the-art at the time, and since then it has made a huge impact on the machine learning community. [Geoffrey Hinton gave an invited talk](#) on this and more recent research Thursday, December 1, at 2:30 pm.

Accessing NeurIPS content:

The tutorials and Invited talks are available on the virtual site to anyone with a NeurIPS account, no NeurIPS registration is required.

All virtual content will be available to the public online 30 days after the conference.

###