



Annual NeurIPS Conference Announces Best Paper Awards for Main Track and Datasets & Benchmark Track

Conference focuses on the research and science in the field of artificial intelligence (AI) and machine learning (ML)

Vancouver, British Columbia, December 11, 2024 — The 38th annual conference on Neural Information Processing Systems ([NeurIPS](#)), a premier conference in artificial intelligence (AI) and machine learning (ML), today announced two Best Main Track Paper Awards, two Main Track Runners-Up, and one Best Datasets and Benchmark Track Paper Award. NeurIPS is being held December 10-15, 2024 at the Vancouver Convention Center.

NeurIPS aims to foster the exchange of research advancements in AI and ML, principally by hosting an annual interdisciplinary academic conference with the highest ethical standards for a diverse and inclusive community. This year's conference has 16,500 registered participants, 4037 main track papers, 460 datasets & benchmark papers, [69 oral sessions](#), [56 workshops](#), [16 Competitions](#), [14 Tutorials](#), [11 Affinity Events](#), [8 Socials](#), and [7 Invited Talks](#).

The search committee for the "Best Paper Award" was nominated by the Program Chairs and the respective Track Chairs, who selected leading researchers with a diverse perspective on machine learning topics. These nominations were approved by the General Chairs and the Diversity, Inclusion and Accessibility (DIA) Chairs.

The best paper award committees were tasked with selecting a handful of highly impactful papers from both tracks of the conference. The search committees considered all accepted NeurIPS papers equally, and made decisions independently based on the scientific merit of the papers, without making separate considerations on authorship or other factors, in keeping with the NeurIPS blind review process.

The best paper and runner up paper awards go to five ground-breaking papers (four main track and one datasets and benchmarks track) that highlight, respectively, a new autoregressive model for vision, new avenues for supervised learning using higher-order

derivatives, improved training of LLMs and inference methods for text2image diffusion and a novel diverse benchmark dataset for LLM alignment.

Best Paper Awards for the main track:

[Visual Autoregressive Modeling: Scalable Image Generation via Next-Scale Prediction](#)

Authors: Keyu Tian, Yi Jiang, Zehuan Yuan, BINGYUE PENG, Liwei Wang

Abstract: This paper introduces a novel visual autoregressive (VAR) model that iteratively predicts the image at a next higher resolution, rather than a different patch in the image following an arbitrary ordering. The VAR model shows strong results in image generation, outperforming existing autoregressive models in efficiency and achieving competitive results with diffusion-based methods. At the core of this contribution lies an innovative multiscale VQ-VAE implementation. The overall quality of the paper presentation, experimental validation and insights (scaling laws) give compelling reasons to experiment with this model.

[Stochastic Taylor Derivative Estimator: Efficient amortization for arbitrary differential operators](#)

Authors: Zekun Shi, Zheyuan Hu, Min Lin, Kenji Kawaguchi

Abstract: This paper proposes a tractable approach to train neural networks (NN) using supervision that incorporates higher-order derivatives. Such problems arise when training physics-informed NN to fit certain PDEs. Naive application of automatic differentiation rules are both inefficient and intractable in practice for higher orders k and high dimensions d . While these costs can be mitigated independently (e.g. for large k but small d , or large d but small k using subsampling) this paper proposes a method, stochastic taylor derivative estimator (STDE) that can address both. This work opens up possibilities in scientific applications of NN and more generally in supervised training of NN using higher-order derivatives.

Runners-up for the main track:

[Not All Tokens Are What You Need for Pretraining](#)

Authors: Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu, yelong shen, Ruochen Xu, Chen Lin, Yujiu Yang, Jian Jiao, Nan Duan, Weizhu Chen

Abstract: This paper presents a simple method to filter pre-training data when training large language models (LLM). The method builds on the availability of a high-quality reference dataset on which a reference language model is trained. That model is then used to assign

a quality score for tokens that come from a larger pre-training corpus. Tokens whose scores have the highest rank are then used to guide the final LLM training, while the others are discarded. This ensures that the final LLM is trained on a higher quality dataset that is well aligned with the reference dataset.

[Guiding a Diffusion Model with a Bad Version of Itself](#)

Authors: Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, Samuli Laine

Abstract: This paper proposes an alternative to classifier free guidance (CFG) in the context of text-2-image (T2I) models. CFG is a guidance technique (a correction in diffusion trajectories) that is extensively used by practitioners to obtain better prompt alignment and higher-quality images. However, because CFG uses an unconditional term that is independent from the text prompt, CFG has been empirically observed to reduce diversity of image generation. The paper proposes to replace CFG by Autoguidance, which uses a noisier, less well-trained T2I diffusion model. This change leads to notable improvements in diversity and image quality.

Best Paper Award for Datasets & Benchmarks track:

[The PRISM Alignment Dataset: What Participatory, Representative and Individualised Human Feedback Reveals About the Subjective and Multicultural Alignment of Large Language Models](#)

Authors: Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Michael Bean, Katerina Margatina, Rafael Mosquera, Juan Manuel Ciro, Max Bartolo, Adina Williams, He He, Bertie Vidgen, Scott A. Hale

Abstract: Alignment of LLMs with human feedback is one of the most impactful research areas of today, with key challenges such as confounding by different preferences, values, or beliefs. This paper introduces the PRISM dataset providing a unique perspective on human interactions with LLMs. The authors collected data from 75 countries with diverse demographics and sourced both subjective and multicultural perspectives benchmarking over 20 current state of the art models. The paper has high societal value and enables research on pluralism and disagreements in RLHF.

Best Paper Award committee for main track: Marco Cuturi (Committee Lead), Zeynep Akata, Kim Branson, Shakir Mohamed, Remi Munos, Jie Tang, Richard Zemel, Luke Zettlemoyer

Best Paper Award committee for dataset and benchmark track: Yulia Gel, Ludwig Schmidt, Elena Simperl, Joaquin Vanschoren, Xing Xie.

About the conference of Neural Information Processing Systems (NeurIPS)

The conference is organized by the Neural Information Processing Systems Foundation, a non-profit corporation whose purpose is to foster insights into solving difficult problems by bringing together researchers from biological, psychological, technological, mathematical and theoretical areas of science and engineering. For more information, please visit the NeurIPS 2024 [website](#) and monitor the [NeurIPS blog](#) for event updates.

Press Contacts:

Jill Miley

Interprose for NeurIPS 2024

Jill.Miley@InterprosePR.com

NeurIPS Communication Chairs

press@neurips.cc