



Top AI Conference ICLR Unveils First Test of Time Award and Best Paper Award Winners

5 Paper Awards and 11 Honorable Mentions receive prestigious recognition and for lasting impact the Test of Time awards debut

Vienna, Austria, 7 May 2024 - The International Conference on Learning Representations (ICLR), the premier gathering of professionals dedicated to the advancement of the many branches of artificial intelligence (AI) and deep learning—has announced for the first time the prestigious Test of Time awards during the twelfth annual ICLR event at [Messe Wien Exhibition and Congress Center](#) in Vienna, Austria. This hybrid event will be live-streamed in the CEST timezone, 7-11 May. The conference [schedule](#) will feature 7 invited keynote speakers, 86 Oral presentations across 32 Oral sessions, and 8 poster sessions where all paper awards will be presented by the authors and discussed. There are over 5,990 participants from 79 countries taking part in ICLR 2024.

For the first time in ICLR's history, the organizers have identified two papers that have withstood the "Test of Time" and made a lasting impact since 2013 and 2014. These winners were announced live on stage during the opening ceremony. To learn more attend the [Test of Time](#) presentations on 8 May in Hall A 8-9 at 2:15 pm CEST.

Winner of the Test of Time Award

[Diederik Kingma](#), [Max Welling](#)
[Auto-Encoding Variational Bayes](#)

Probabilistic modeling is one of the most fundamental ways in which we reason about the world. This paper spearheaded the integration of deep learning with scalable probabilistic inference, empowering a new class of generative models that enables both realistic generation of images as well as inferring scientific properties from partial or corrupted measurements.

Test of Time Runner-Up

[Christian Szegedy](#), [Wojciech Zaremba](#), [Ilya Sutskever](#), [Joan Bruna](#), [Dumitru Erhan](#), [Ian Goodfellow](#), [Rob Fergus](#)
[Intriguing properties of neural networks](#)

With the rising popularity of deep neural networks in real applications, it is important to understand when and how neural networks might behave in undesirable ways. This paper highlighted the issue that neural networks can be vulnerable to small almost imperceptible variations to the input. This idea helped spawn the area of adversarial attacks (trying to fool a neural network) as well as adversarial defense (training a neural network to not be fooled).

Outstanding Paper Awards

Each year, ICLR organizers receive thousands of paper submissions. This year was no exception, with a 47% increase in paper submissions from 2023. Out of the 7,262 research papers submitted 2,260 (31.1%) were accepted and invited to present during the Oral and Poster sessions. Peer reviews were completed by 60 Senior Area Chairs, 624 Area Chairs, and 8,950 Reviewers, with each paper receiving at least 3 unique reviews.



The Awards Committee includes Eunsol Choi, Katja Hofmann, Ming-Yu Liu, Nan Jiang, Stephan Günnemann, Suvrit Sra, Thomas Kipf and Volkan Cevher. In total, there are 5 Outstanding Paper winners and 11 Honorable Mentions.

Outstanding Award Winners

[Generalization in diffusion models arises from geometry-adaptive harmonic representations](#)

Zahra Kadkhodaie, Florentin Guth, Eero P. Simoncelli, Stéphane Mallat

[Learning Interactive Real-World Simulators](#)

Sherry Yang, Yilun Du, Seyed Kamyar Seyed Ghasemipour, Jonathan Tompson, Leslie Pack Kaelbling, Dale Schuurmans, Pieter Abbeel

[Never Train from Scratch: Fair Comparison of Long-Sequence Models Requires Data-Driven Priors](#)

Ido Amos, Jonathan Berant, Ankit Gupta

[Protein Discovery with Discrete Walk-Jump Sampling](#)

Nathan C. Frey, Dan Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, Andreas Loukas, Vladimir Gligorijevic, Saeed Saremi

[Vision Transformers Need Registers](#)

Timothée Darcet, Maxime Oquab, Julien Mairal, Piotr Bojanowski

Honorable Mentions

[Amortizing intractable inference in large language models](#)

Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, Nikolay Malkin

[Approximating Nash Equilibria in Normal-Form Games via Stochastic Optimization](#)

Ian Gemp, Luke Marris, Georgios Piliouras

[Beyond Weisfeiler-Lehman: A Quantitative Framework for GNN Expressiveness](#)

Bohang Zhang, Jingchu Gai, Yiheng Du, Qiwei Ye, Di He, Liwei Wang

[Flow Matching on General Geometries](#)

Ricky T. Q. Chen, Yaron Lipman

[Is ImageNet worth 1 video? Learning strong image encoders from 1 long unlabelled video](#)

Shashanka Venkataramanan, Mamshad Nayeem Rizve, Joao Carreira, Yuki M. Asano, Yannis Avrithis



[Meta Continual Learning Revisited: Implicitly Enhancing Online Hessian Approximation via Variance Reduction](#)

Yichen Wu, Long-Kai Huang, Renzhen Wang, Deyu Meng, Ying Wei

[Model Tells You What to Discard: Adaptive KV Cache Compression for LLMs](#)

Suyu Ge, Yunan Zhang, Liyuan Liu, Minjia Zhang, Jiawei Han, Jianfeng Gao

[Proving Test Set Contamination in Black-Box Language Models](#)

Yonatan Oren, Nicole Meister, Niladri S. Chatterji, Faisal Ladhak, Tatsunori Hashimoto

[Robust agents learn causal world models](#)

Jonathan Richens, Tom Everitt

[The mechanistic basis of data dependence and abrupt learning in an in-context classification task](#)

Gautam Reddy

[Towards a statistical theory of data selection under weak supervision](#)

Germain Kolossov, Andrea Montanari, Pulkit Tandon

Invited Talks

Kyunghyun Cho [*Machine Learning in Prescient Design's Lab-in-the-Loop Antibody Design*](#)

Abstract: Together with two other co-founders, Rich Bonneau and Vlad Gligorijevic, I founded Prescient Design in January 2021, in order to build a lab-in-the-loop protein design platform based on our earlier research. Prescient Design was fully acquired by Genentech (Roche) on August 2021, and began to focus more specifically on antibody design. It has been more than three years since its founding and more than 2.5 years since the acquisition. In this talk, I will share Prescient Design's lab-in-the-loop antibody design, both the platform and the outcome, as well as what went behind in building this platform from the perspective of machine learning.

Priya Donti [*Why your work matters for climate in more ways than you think*](#)

Abstract: Climate change is one of the most pressing issues of our time, requiring rapid transformation across virtually every sector of society. In this talk, I describe what this means for research and practice in AI. AI has a multi-faceted relationship with climate change, through a combination of its direct environmental footprint, the impacts of its applications (both good and bad), and the broader systemic shifts it induces. Ultimately, most work in AI has significant implications for climate action, whether or not it is viewed as traditionally “climate-relevant.” Given this, I discuss how the AI community can better align its work with climate action: through the kinds of methods we develop, the kinds of applications we work on, the choices we make while working on these applications, and the ways we communicate with the public about our work.

Kate Downing [*Copyright Fundamentals for AI Researchers*](#)

Abstract: This talk will cover fundamental legal principles all AI researchers should understand about copyright law. This talk will explore the current state of copyright law with respect to AI in the U.S., potential claims and defenses, as well as practical tips for minimizing legal risk.



Raia Hadsell, Google DeepMind, [*Learning through AI's winters and springs: unexpected truths on the road to AGI*](#)

Abstract: After decades of steady progress and occasional setbacks, the field of AI now finds itself at an inflection point. AI products have exploded into the mainstream, we've yet to hit the ceiling of scaling dividends, and the community is asking itself what comes next. In this talk, Raia will draw on her 20 years experience as an AI researcher and AI leader to examine how our assumptions about the path to Artificial General Intelligence (AGI) have evolved over time, and to explore the unexpected truths that have emerged along the way. From reinforcement learning to distributed architectures and the potential of neural networks to revolutionize scientific domains, Raia argues that embracing lessons from the past offers valuable insights for AI's future research roadmap.

Moritz Hartd [*The emerging science of benchmarks*](#)

Abstract: Benchmarks are the keystone that hold the machine learning community together. Growing as a research paradigm since the 1980s, there's much we've done with them, but little we know about them. In this talk, I will trace the rudiments of an emerging science of benchmarks through selected empirical and theoretical observations. Specifically, we'll discuss the role of annotator errors, external validity of model rankings, and the promise of multi-task benchmarks. The results in each case challenge conventional wisdom and underscore the benefits of developing a science of benchmarks.

Devi Parikh [*Stories from my life*](#)

Abstract: This is going to be an unusual AI conference keynote talk. When we talk about why the technological landscape is the way it is, we talk a lot about the macro shifts – the internet, the data, the compute. We don't talk about the micro threads, the individual human stories as much, even though it is these individual human threads that cumulatively lead to the macro phenomenon. We should talk about these stories more! So that we can learn from each other, inspire each other. So we can be more robust; more effective in our endeavors. By strengthening our individual threads and our connections, we can weave a stronger fabric together. This talk is about some of my stories from my 20-year journey so far – about following up on all threads, about learnt reward functions, about fleeting opportunities, about multidimensional impact landscapes, and about curiosity for new experiences. It might seem narcissistic, but hopefully it will also feel authentic and vulnerable. And hopefully you will get something out of it.

Jie Tang [*The ChatGLM's Road to AGI*](#)

Abstract: Large language models have substantially advanced the state of the art in various AI tasks, such as natural language understanding and text generation, and image processing, multimodal modeling. In this talk, we will first introduce the development of AI in the past decades, in particular from the angle of China. We will also talk about the opportunities, challenges, and risks of AGI in the future. In the second part of the talk, we will use ChatGLM, an alternative but open-sourced model to ChatGPT, as an example to explain our understanding and insights derived during the implementation of the model.

The Organizers



“ICLR Vienna will facilitate the gathering of deep learning and artificial intelligence researchers to share and exchange ideas as well as network,” said Been Kim, ICLR 2024 general chair. “We are particularly excited to introduce this year the Test of Time award which is a testament to research conducted over 10 years ago with a lasting impact. The ICLR program chairs invited keynote speakers to address important topics as the AI research community leads efforts to benefit humanity and address challenges.”

The program chairs, Swarat Chaudhuri, Mohammad Emtiyaz Khan, Yizhou Sun and Katerina Fragkiadaki along with the entire organizing committee are proud to present to you ICLR 2024 in Vienna.

Global participants at ICLR span 79 countries with various backgrounds, from academic and industrial researchers to entrepreneurs and engineers, to graduate students and post-doctorates. ICLR continues to pursue inclusivity and efforts to reach a broader audience, employing activities such as [mentoring programs](#) (formally "office hours") and hosting social meetups on a global scale. Explore global, cutting-edge research on all aspects of deep learning used in artificial intelligence, statistics and data science, as well as important application areas such as machine vision, computational biology, speech recognition, text understanding, gaming, and robotics.

All recordings will become available to the public a month after the event ends. To stay informed about ICLR, visit the [website](#) and monitor the [ICLR Blog](#) and [X](#).

About ICLR

The International Conference on Learning Representations (ICLR) is the premier gathering of professionals dedicated to the advancement of the branch of artificial intelligence called representation learning but generally referred to as deep learning. For more information about ICLR visit: <https://iclr.cc/>.

Media Contact:

Becky Obbema

Interprose, PR for ICLR

becky.obbema@interprosepr.com / press@iclr.cc