NeurIPS 2025 – Fact Sheet

39th Annual Conference on Neural Information Processing Systems
San Diego, California, USA & Mexico City, Mexico

Registration Numbers

- Total Registrants: approximately 26382
- In-person:
    - San Diego: approximately 24,500
    - Mexico City: 463
- Virtual: 2366

Past Attendance & Locations

- 2024 — Vancouver, Canada — 19,756 (hybrid)
- 2023 — New Orleans, USA — 16,382
- 2022 — New Orleans, USA — 15,390
- 2021 — Virtual — 17,091
- 2020 — Virtual — 22,823
- 2019 — Vancouver, Canada — 13,000
- 2018 — Montreal, Canada — 8,648
- 2017 — Long Beach, USA — 8,008
- 2016 — Barcelona, Spain — 5,231
- 2015 — Montreal, Canada — 3,852
- 2014 — Montreal, Canada — 2,581
- 2013 — Lake Tahoe, USA — 1,994
- 2012 — Lake Tahoe, USA — 1,676
- 2011 — Granada, Spain — 1,452
- 2010 — Vancouver, Canada — 1,354

Program Overview

- **Tracks: 3**
    - Number of Orals: 87
    - Number of Posters: 5,290
    - Position Paper Track Panels: 2
- **Invited Speakers: 6**
- **Awards: 3**
    - Test of Time
    - Sejnowski-Hinton Prize
    - Best Papers 2025
- **Expo Programming**
    - Demonstrations: 20
    - Talk/Panels: 19
    - Workshops:11
- **Competitions:**
    - San Diego: 18
    - Mexico City: 1
- **Tutorials**
    - San Diego: 14
    - Mexico City: 6
- **Workshops**
    - San Diego: 55
    - Mexico City: 8

Review Stats

- Submissions: 21,575
- Accepted: 5,290
- Review Pool
    - Total Senior Area Chairs: 240
    - Total Area Chairs: 1,985
    - Total technical reviewers: 21,921
    - Total ethics reviewers: 641

Invited Speakers

Kyunghyun Cho

Glen de Vries Professor of Health Statistics, NYU; Executive Director of Frontier Research, Prescient Design, Genentech
Website: kyunghyuncho.me

Cho co-developed the Gated Recurrent Unit (GRU) and has been central to advances in neural machine translation and sequence-to-sequence learning. He is a CIFAR Fellow in Learning in Machines & Brains, winner of the 2021 Samsung Ho-Am Prize in Engineering, and has served as Program Chair for ICLR 2020, NeurIPS 2022, and ICML 2022.

Yejin Choi

Professor of Computer Science, Stanford University; Dieter Schwarz Foundation Senior Fellow, Stanford HAI; Distinguished Scientist, NVIDIA
Website: yejinc.github.io

Choi works at the intersection of language understanding and commonsense reasoning. A MacArthur Fellow (2022) and Time Magazine's Most Influential People in AI (2023), she has earned multiple Test of Time Awards and Best Paper Awards across ACL, EMNLP, ICML, and NeurIPS.

Melanie Mitchell

Professor, Santa Fe Institute
Website: melaniemitchell.me

Mitchell studies AI, cognitive science, analogy-making, and complex systems. Author of *Artificial Intelligence: A Guide for Thinking Humans* and *Complexity: A Guided Tour* (winner of the 2010 Phi Beta Kappa Science Book Award). She received her PhD under Douglas Hofstadter and co-developed the Copycat cognitive architecture.

Andrew Saxe

Professor of Theoretical Neuroscience & Machine Learning, Gatsby Computational Neuroscience Unit and Sainsbury Wellcome Centre, UCL
Website: saxelab.org

Saxe develops mathematical theories of deep learning, including exact solutions for learning dynamics in deep linear networks and formal connections between artificial and biological learning. CIFAR Fellow in Learning in Machines & Brains; winner of the 2019 Wellcome Trust Beit Prize.

Richard Sutton

Research Scientist, Keen Technologies; Professor, University of Alberta; Chief Scientific Advisor, Amii; Chief Scientific Officer, ExperienceFlow.ai
Website: incompleteideas.net

Co-developer of temporal difference learning and policy gradient methods, Sutton won the 2024 Turing Award (with Andrew Barto). Co-author of *Reinforcement Learning: An Introduction*. His research focuses on computational principles of learning and decision-making.

Zeynep Tufekci

Henry G. Bryant Professor of Sociology and Public Affairs, Princeton University; New York Times Columnist
Website: zeynep.me

Tufekci examines the relationships between technology, society, and complex social systems. A 2022 Pulitzer Prize finalist, she is known for *Twitter and Tear Gas* and for her writing on digital public spheres. She is also affiliated with the Berkman Klein Center at Harvard.

Affinity Groups Presenting

San Diego

- Indigenous in AI/ML
- LatinX in AI
- Muslims in ML
- New in ML
- Neurodivergent in AI
- Queer in AI
- Women in Machine Learning
- Black in AI
- Global South in AI

Mexico City

- LatinX in AI (CDMX)
- Queer in AI (CDMX)
- Women in Machine Learning (CDMX)
- Joint Affinity Social (LXAI + Queer in AI + WiML)

Awards (2025)

Sejnowski–Hinton Prize

Winner: *Random synaptic feedback weights support error backpropagation for deep learning* (2016) Timothy Lillicrap, Daniel Cownden, Douglas Tweed, Colin Akerman

Impact: Introduced feedback alignment, the first biologically plausible solution to the weight transport problem.

Test of Time Award

Winner: *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks* Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun (NeurIPS 2015)

Impact: Defined modern two-stage object detection; introduced the RPN; >56,700 citations; memorial recognition for Dr. Jian Sun.

Best Paper Awards

Seven groundbreaking papers, including four best papers (one of which is from the datasets and benchmarks track) and three runner-ups were chosen this year. The seven papers highlight advances in diffusion model theory, self-supervised reinforcement learning, attention mechanisms for large language models, reasoning capabilities in LLMs, online learning theory, neural scaling laws, and benchmarking methodologies for language model diversity. You can read more about the papers and their selection in our blog post.

Winners

- **Artificial Hivemind: The Open-Ended Homogeneity of Language Models (and Beyond)**
  Authors: Liwei Jiang, Yuanjun Chai, Margaret Li, Mickel Liu, Raymond Fok, Nouha Dziri, Yulia Tsvetkov, Maarten Sap, Yejin Choi

- **Gated Attention for Large Language Models: Non-linearity, Sparsity, and Attention-Sink-Free**
  Authors: Zihan Qiu, Zekun Wang, Bo Zheng, Zeyu Huang, Kaiyue Wen, Songlin Yang, Rui Men, Le Yu, Fei Huang, Suozhi Huang, Dayiheng Liu, Jingren Zhou, Junyang Lin

- **1000 Layer Networks for Self-Supervised RL: Scaling Depth Can Enable New Goal-Reaching Capabilities**
  Authors: Kevin Wang, Ishaan Javali, Michał Bortkiewicz, Tomasz Trzcinski, Benjamin Eysenbach

- **Why Diffusion Models Don't Memorize: The Role of Implicit Dynamical Regularization in Training**
  Authors: Tony Bonnaire, Raphaël Urfin, Giulio Biroli, Marc Mezard

---

Runners-Up

- **Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?**
  Authors: Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, Gao Huang

- **Optimal Mistake Bounds for Transductive Online Learning**
  Authors: Zachary Chase, Steve Hanneke, Shay Moran, Jonathan Shafer

- **Superposition Yields Robust Neural Scaling**
  Authors: Yizhou Liu, Ziming Liu, Jeff Gore

About the Conference

The conference was founded in 1987 and is now a multi-track interdisciplinary annual meeting that includes invited talks, demonstrations, symposia, and oral and poster presentations of refereed papers. Along with the conference is a professional exposition focusing on machine learning in practice, a series of tutorials, and topical workshops that provide a less formal setting for the exchange of ideas.

More about the Neural Information Processing Systems foundation