Workshop on Scaling Post-training for LLMs (SPOT)

ICLR 2026 Workshop Proposal

Website: https://spoticlr.github.io/

---

1 Workshop Summary

Post-training, encompassing techniques like Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL), is no longer a mere final step for task-specific adaptation. It is evolving into a compute-intensive phase in its own right, crucial for unlocking the full potential of foundational models and optimizing for critical downstream behaviors. Yet, the science of post-training, at scale, remains in its infancy.

Both the academic and industrial research landscape is already signaling the importance of this shift. While models from organizations like DeepSeek and xAI (with Grok) have reportedly dedicated post-training compute budgets that are a significant fraction of their pre-training compute, parallely academic labs are focusing on developing principled ways to scale post-training [Tan et al., 2025, He and Cao, 2025, Du et al., 2025], and democratizing frameworks like SkyRL [Griggs et al., 2025]. This dramatic increase in RL compute is only growing across frontier LLM generations, with more than 10× increase from o1 to o3 [OpenAI, 2025] and a similar leap from Grok-3 to Grok-4 [xAI Team, 2025]. This demonstrates that scaling post-training is both a future-looking academic question as well as a present-day industrial imperative.

This **Workshop on Scaling Post-Training for LLMs (SPOT)** is motivated by the urgent need to **establish rigorous and scalable methodologies, design choices, and approaches for post-training**. While today's design choices in pre-training are made with a core focus on their ability to scale, a similar "scaling laws" mindset for post-training is largely absent. Our goal is to catalyze a systematic understanding of how post-training scales, across algorithms, data regimes, infrastructure, and objectives, and to identify the open questions that must be addressed to turn post-training into a science of its own.

While many existing workshops focus on specific applications of post-training, such as reasoning, planning, or agentic behavior, our goal is distinct. This workshop centers on the foundational principles that make post-training itself scale effectively, regardless of its end use. We aim to examine the design choices and bottlenecks that influence post-training efficiency, stability, and asymptotic performance across all paradigms - SFT, RL, and beyond.

Overview of Discussion Topics– The workshop will cover a range of topics, including, but not limited to:

- **Algorithmic insights to scaling Post-training effectively:** Which SFT and RL formulations, objective choices, and reward-modeling strategies enable stable and compute-efficient scaling? Emerging studies have begun probing off-policy training [Zheng et al., 2025, Piche et al., 2025], the relative learning contributions of SFT and RL [Schulman and Lab, 2025], and the role of diversity in reward signals [Zhu et al., 2025, Chen et al., 2025]. Each of these factors impacts both the asymptotic performance attainable at scale and the sample efficiency required to reach it. An ideal goal is to understand how such algorithmic insights can be unified into a scaling-law perspective for post-training.
- **Specialized systems to accelerate Post-training:** Large-scale post-training introduces unique engineering bottlenecks, often distinct from pre-training. Topics such as training–inference kernel mismatches [He and

Lab, 2025], rollout latency [Fu et al., 2025], distributed training and inference [Mika et al., 2025, Intellect, 2025], or KV-cache management under long contexts [Aghajohari et al., 2025] directly affect scalability and throughput, and hence efficiency of post-training.

- **Architectural choices and their impact at scale:** Architectural advances, such as local-global attention [Team et al., 2025, Meta AI, 2025], and sparse mixture-of-experts (MoE) [Fedus et al., 2022, Obando-Ceron et al., 2024], have redefined the compute–capability tradeoff in post-training. Similar fundamental questions remain: Is a single, unified post-trained model the optimal path, or should we focus on developing ecosystems of specialized models that can be orchestrated as agents? How do architectural choices influence sample efficiency, generalization, and stability at scale?
- **Role of Data in scaling Post-training:** Data remains a central driver of performance and efficiency at scale. In pre-training, it's known to affect both the performance as well as efficiency [Li et al., 2025, Everett, 2025]. We need to understand data's corresponding principles in post-training. Moreover, as the scale of post-training continues to grow, new questions emerge around the role of synthetic data. What are the scaling laws of synthetic data in post-training? How does it affect the diversity and overall generalization?
- **Environment design and safety in real-world Post-training:** Post-training expands beyond code and mathematics into multimodal, embodied, and experimental domains that require lab-based or physically grounded setups (e.g., RL for scientific discovery). Each introduces distinct challenges for environment design, evaluation fidelity, and safety.
- **Efficient and Dense Feedback from the environment:** As models grow larger and tasks more complex, verifying post-training outcomes has become a major bottleneck. In many cases, evaluation cost rivals or exceeds training cost [Zhang et al., 2025]. Building efficient and dense verification methods remains central to scalable post-training.

This workshop aims to bring together diverse perspectives from academic and industrial researchers and practitioners, to share practical experiences, and to outline a clear research direction toward building a principled science of post-training at scale.

2 Invited Speaker/Panelists Information

This workshop will feature senior scholars with extensive expertise in LLMs, RL, reasoning, training infrastructure, data curation, who will deliver keynote presentations about science of scaling post-training. The invited speakers are as follows:

Speaker/Panelist	Affiliation	Tentative Topic/Area
Nouha Dziri	Allen Institute of AI	Safety and general post-training
Rishabh Agarwal	Periodic Labs	Scaling Laws for RL
Aakanksha Chowdhary	Stanford / Reflection AI	Open source RL at large scale
Sanmi Koyejo	Stanford	Trustworthiness for large scale post-trained models
Jiantao Jiao	UC Berkeley / Nvidia	Large scale post-training for verifiable agents
Junyang Lin	Qwen Team	Open source post-training at large scale
Sami Jaghouar	Prime Intellect	Large scale distributed post-training
Yuandong Tian	Meta	Scalable post-training
Cho-Jui Hsieh	UCLA / Google	Effect of data on post-training

Table 1: List of speakers and panelists, their affiliations, and tentative topic of their talks.

- **Nouha Dziri, Allen Institute of AI (Confirmed)**, Nouha Dziri is a Research Scientist at the Allen Institute for AI (AI2), where she co-leads the safety and post-training efforts for building OLMo [Groeneveld et al., 2024], a highly capable and truly open large language model (LLM) designed to advance AI research and democratize access to powerful language technologies. Her research focuses on understanding the limits of Transformer architectures and their inner mechanisms. Her work “Faith and Fate” explores the boundaries of Transformers on reasoning tasks. She also studies alignment in LLMs, contributing to research such as “Roadmap to Pluralistic Alignment,” “Finegrained RLHF,” and “RewardBench,” the first comprehensive evaluation framework for reward models. Her research has been featured in TechCrunch, Le Monde, The Economist, Science News, and Quanta Magazine. Previously, she collaborated with several leading researchers in the field, including Siva Reddy at Mila/McGill, Hannah Rashkin, Tal Linzen, David Reitter, Diyi Yang, and Tom Kwiatkowski at Google Research NYC, as well as Alessandro Sordoni and Geoff Gordon at Microsoft Research Montreal.

- **Rishabh Agarwal, Periodic Labs (Confirmed)**, is a founding member of *Periodic Labs* and an Adjunct Professor at McGill University. His research focuses on the intersection of reinforcement learning (RL) and large language models (LLMs), aiming to build more general and capable learning systems. Previously, he worked on reinforcement learning and reasoning at Meta, and before that was a Staff Research Scientist at Google DeepMind, contributing to advances in generalization and sample-efficient learning. Rishabh completed his PhD at Mila under the supervision of Aaron Courville and Marc Bellemare, and also spent a year with Geoffrey Hinton’s team at Google Brain, Toronto. He earned his undergraduate degree in Computer Science and Engineering from IIT Bombay, and his prior work on Deep Reinforcement Learning has been recognized with an Outstanding Paper Award at NeurIPS.
- **Aakanksha Chowdhary, Stanford / Reflection AI (Tentative)**, is a founding member of Reflection AI where she is advancing the frontier of agentic large language models by applying reinforcement learning techniques to develop autonomous, self-improving AI agents. She also holds the position of Adjunct Professor at Stanford University, she co-teaches CS329A: Self-Improving AI Agents (Fall/Winter 2025) and serves as Program Chair for MLSys 2026. Previously, she was the technical lead of Google’s 540-billion-parameter PaLM model and a lead researcher on the Gemini project, contributing to large-scale pre-training, scaling, and fine-tuning of large language models. Before Google, she served as a technical lead for several interdisciplinary research programs at Microsoft Research and Princeton University spanning machine learning and distributed systems. She earned her Ph.D. in Electrical Engineering from Stanford University, where her dissertation received the Paul Baran Marconi Young Scholar Award for outstanding scientific contributions in communications and the Internet. Her work has been recognized with Outstanding Paper Awards at MLSys 2022 and 2023, among other honors.
- **Sanmi Koyejo, Stanford University (Confirmed)**, is an Assistant Professor in the Department of Computer Science at Stanford University. Koyejo was previously an Associate Professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. Koyejo’s research interests are in developing the principles and practice of trustworthy machine learning, focusing on applications to neuroscience and healthcare. Koyejo completed a Ph.D. at the University of Texas at Austin, and postdoctoral research at Stanford University. Koyejo has been the recipient of several awards, including a best paper award from the conference on uncertainty in artificial intelligence, a Skip Ellis Early Career Award, a Sloan Fellowship, a Terman faculty fellowship, an NSF CAREER award, a Kavli Fellowship, an IJCAI early career spotlight, and a trainee award from the Organization for Human Brain Mapping. Koyejo spends time at Google as a part of the Brain team, serves on the Neural Information Processing Systems Foundation Board, the Association for Health Learning and Inference Board, and as president of the Black in AI organization.
- **Jiantao Jiao, UC Berkeley / NVIDIA (Confirmed)**, is a professor of Electrical Engineering & Computer Sciences and Statistics at UC Berkeley, where his work spans generative AI and foundation models. He is deeply involved in fostering collaboration between academia and industry, with research funded by the NSF, OpenAI, Meta, Google, among others. He also co-founded and served as CEO of Nexusflow.ai, a startup focused on the post-training of LLM agents, which was later acquired by NVIDIA. He is currently Director of Research & Distinguished Scientist at NVIDIA, where he is focused on advancing the frontier of Artificial Superintelligence.
- **Junyang Lin, Qwen Team (Tentative)**, is a researcher and engineer at the M6 Team, DAMO Academy, Alibaba Group, and a core maintainer of the Qwen Team, where he builds large language and multimodal models. His work focuses on large-scale pretraining, multimodal learning, and downstream applications. He led the development of M6, China’s first large-scale multimodal pretraining system, and OFA, a unified

multimodal multitask framework achieving state-of-the-art results on benchmarks such as MSCOCO and VQA. His research bridges large-scale model training with real-world applications across Alibaba's AI ecosystem.

- **Sami Jaghouar, Prime Intellect (Confirmed)**, is a founding research engineer at Prime Intellect, where he leads work on decentralized and large-scale model training. He is one of the core authors of INTELLECT-1, a globally distributed 10B parameter LLM, and contributed to INTELLECT-2, a 32B reasoning model trained via asynchronous RL. His interests lie at the intersection of distributed learning, reinforcement learning, and infrastructure for scale, with a vision toward open and community-driven AI.
- **Yuandong Tian, Meta (Confirmed)**, is a Research Scientist Director in Meta AI Research (FAIR), leading the group of reasoning, planning and decision-making with Large Language Models (LLMs). He is the project lead for OpenGo project that beats professional players with a single GPU during inference, serves as the main mentor of StreamingLLM and GaLore that improve the training and inference of LLM, and is the first-author recipient of 2021 ICML Outstanding Paper Honorable Mentions DirectPred and 2013 ICCV Marr Prize Honorable Mentions Hierarchical Data Driven Descent, and also received the 2022 CGO Distinguished Paper Award CompilerGym. Prior to that, he worked in Google Self-driving Car team in 2013-2014 and received a Ph.D in Robotics Institute, Carnegie Mellon University in 2013.
- **Cho-Jui Hsieh, UCLA / Google (Confirmed)**, is an Associate Professor of Computer Science at UCLA and a Research Scientist at Google. His research focuses on algorithms and optimization techniques for large-scale machine learning, including improving model efficiency, training speed, prediction performance, and robustness. He received his Ph.D. in Computer Science from the University of Texas at Austin, advised by Prof. Inderjit Dhillon, and his M.S. from National Taiwan University under Prof. Chih-Jen Lin. Before joining UCLA, he was an Assistant Professor at UC Davis and a visiting scholar at Google.

3 Tentative Schedule and Organizational Details

3.1 Tentative Schedule

Our workshop will be conducted in an in-person format. The program will feature six invited talks, each lasting 35 minutes (30 minutes for the presentation and 5 minutes for Q&A), along with six contributed talks of 10 minutes each, selected from submitted papers. In addition, there will be two ∼1-hour poster sessions and a 1-hour panel discussion focused on the future of scaling post-training. Although the workshop will be in person, to foster engagement and interaction among remote attendees, who could not join in-person, a dedicated chat platform will be available, ensuring seamless communication throughout the workshop.

Table 2: Workshop Schedule

Time	**Arrangement**
08:50–09:00	Introduction and Opening Remarks
09:00–09:35	Invited Talk #1
09:35–10:10	Invited Talk #2
10:10–10:40	Oral Presentations (3 presentations, 10 mins each)
10:40–11:00	Coffee Break
11:00–11:35	Invited Talk #3
11:35–12:30	Poster Session 1
12:30–13:30	*Break (Lunch)*
13:30–14:05	Invited Talk #4
14:05–14:40	Invited Talk #5
14:40–15:10	Oral Presentations (3 presentations, 10 mins each)
15:10–15:45	Invited Talk #6
15:45–16:45	Panel Discussion: Prospects & Pitfalls of Scaling Post-Training of LLMs
16:45–17:45	Poster Session 2
17:45–18:00	Paper Award and Closing Remarks

3.2 Estimated Audience Size

Our workshop is designed for researchers and practitioners with a strong interest in machine learning and NLP, with a focus on LLMs Post-training. Based on the success of previous workshops on related topics, the relevance of the proposed agenda, and the appeal of our invited speakers, we anticipate a consistent in-person attendance of 300-400 participants, with a total reach of 500-600 attendees over the course of the event.

3.3 Important Dates for Review Process

We will follow the suggested dates by ICLR.

- Workshop paper submission deadline: January 30, 2026.
- Workshop paper notification date: February 27, 2026.
- Final workshop program, camera-ready: April 1, 2026.

3.4 Support for Tiny or Short Papers

In addition to full-length papers (6–8 pages), we will follow the ICLR workshop guidelines by offering a Tiny/Short Papers track. This track will include 2–4 page submissions focused on implementations and evaluations of unpublished but simple ideas, moderate yet self-contained theoretical results, follow-up experiments or re-analyses of published work, or fresh perspectives on existing publications. This format provides an opportunity for underrepresented, resource-constrained, or early-career researchers to receive feedback and engage in discussions at ICLR.

3.5 Funding

We are currently in discussion with Liner AI regarding a sponsorship of approximately $5,000 USD. This support will be directed toward enabling the participation of early-career researchers and practitioners, with a focus on individuals from underrepresented backgrounds. The funds will primarily be used to cover ICLR registration fees and travel expenses for the aforementioned researchers.

4 Diversity and Inclusion Statement

The organizing committee is strongly committed to cultivating diversity and inclusion across research disciplines, career stages, and personal identities. From the outset, we have prioritized balanced representation among our invited speakers and contributors. To further support inclusion, we will actively engage with established communities and networks, such as WiML, Black in AI, LXAI, and Queer in AI, to encourage participation from women and other underrepresented groups through paper submissions and open discussions. Key aspects of our workshop's diversity and inclusion plan are outlined below:

- **Diversity of Topics:** The workshop encompasses a broad spectrum of themes related to post-training, including theoretical foundations, algorithmic design, system and infrastructure optimization, data-centric analysis, and the study of post-training across diverse environments beyond reasoning or coding tasks. We also highlight the safety, reliability, and alignment dimensions of post-training. This diversity of focus areas is intended to bring together researchers from multiple subcommunities, foster interdisciplinary dialogue, and catalyze new collaborations. Our invited speakers similarly reflect this topical breadth.

- **Diversity of Speakers:** Our speaker lineup reflects diversity across gender, ethnicity, career stage, and professional background. Two of our invited speakers identify as female, and others identify as male. The group includes individuals identifying as Black, Asian, Indian, European, Mediterranean, and American. In terms of seniority, we have two assistant professors, one associate professor, three individuals with joint appointments in academia and industry, five primarily from industry, and three from early-stage startups. Each speaker brings distinct expertise from a different area of post-training research and application.

- **Diversity of Organizers:** Our organizing committee consists of eight members, including two who identify as female. The team spans multiple career stages—from faculty (one assistant professor and one full professor) to graduate students (one second-year and two senior PhD students), as well as professionals in industry at varying levels, including an early-career ML engineer at Microsoft, a software engineer at Meta, a research scientist, and a principal scientist at GDM. Organizers represent diverse institutions such as UT Austin, UC Berkeley, UCL, Microsoft, Meta, and Google. Notably, four of the eight organizers have extensive prior experience organizing workshops at major venues including ICLR, NeurIPS, ICCV and ICML.

5 Organizing Committee

The organizing committee consists of both junior and senior members working broadly along scalable post-training recipes. Several members have extensive experience organizing workshops at major machine learning conferences, including NeurIPS, ICLR, ICCV and ICML.

- **Devvrit Khatri (PhD student at UT Austin)**: Devvrit is a 5^{th} year PhD student at UT Austin, advised by Prof. Inderjit Dhillon. His research interests lie in Large Scale Machine Learning, with a focus on efficiency and scalability. His work encompasses building fundamental ML algorithms that exhibit strong empirical performance coupled with real-world deployability. Many of his works, like MatFormer [Devvrit et al., 2024], and DiskANN [Jayaram Subramanya et al., 2019] have been integrated into the workflows of Google and Microsoft, among other organizations.
- **Sewon Min (Assistant Professor at UC Berkeley)**. Sewon is an Assistant Professor at the University of California, Berkeley (EECS) and a Research Scientist at the Allen Institute for AI (AI2). She is affiliated with the Berkeley Artificial Intelligence Research (BAIR) Lab and the Berkeley NLP Group. Her research focuses on Natural Language Processing and Machine Learning, particularly on Large Language Models (LLMs), with an emphasis on understanding and advancing how LLMs leverage large-scale text corpora. She received her Ph.D. in Computer Science & Engineering from the University of Washington, advised by Hannaneh Hajishirzi and Luke Zettlemoyer. During her Ph.D., she was a part-time visiting researcher at Meta AI/FAIR, and interned at Google Research and Salesforce Research. She holds a B.S. in Computer Science and Engineering from Seoul National University.
- **Rishabh Tiwari (PhD Student at UC Berkeley)**: Rishabh is a Ph.D. student in the Berkeley Artificial Intelligence Research (BAIR) lab at UC Berkeley's Department of Electrical Engineering and Computer Sciences (EECS), advised by Prof. Kurt Keutzer. His research focuses on efficient deep learning, with recent works on mitigating simplicity bias without bias labels and advancing network architecture optimization through network pruning algorithms. Before joining Berkeley, Rishabh was a Pre-Doctoral Researcher at Google DeepMind, where they worked with Dr. Pradeep Shenoy on developing algorithms to improve the group-robustness of neural networks. Beyond academia, Rishabh serves as a founding member and senior researcher at Transmute AI Labs, a nonprofit research organization dedicated to mentoring undergraduate students and fostering independent, high-impact AI research.
- **Nan Rosemary Ke (Research Scientist at Google DeepMind)**: Rosemary is a Research Scientist at Google DeepMind, where she builds reasoning systems for math, code, and decision-making. Her recent work focuses on advancing the reasoning capabilities of the Gemini family, including Gemini and Gemini 2.5, where she has been a key contributor to reasoning research. Previously, her research centered on causality and modularity in deep learning, and her current interests broadly span strengthening the reasoning capabilities of large models across diverse domains. She completed her Ph.D. at Mila with Professors Yoshua Bengio and Chris Pal. During her Ph.D., she also spent time at Google DeepMind, Meta AI Research, and Microsoft Research Montreal. She was named a Rising Star in Machine Learning and a Rising Star in EECS, and received the Facebook Fellowship.
- **Gagan Jain (Applied Scientist at Microsoft AI)**: Gagan is an Applied Scientist at Microsoft AI. His research focuses on generative modeling and efficiency, with current explorations in retrieval augmented generation, planned discrete diffusion and game-theoretic perspectives on post-training. Previously, he was a pre-doctoral researcher in the Machine Learning and Optimization team at Google DeepMind, Bangalore, advised by Dr. Sujoy Paul and Dr. Prateek Jain. At DeepMind, he worked on Gemini pre-training, efficient large multi-modal architectures, and theoretical guarantees for collaborative bandit approaches. Gagan earned his B.Tech (Hons.) in Mechanical Engineering from IIT Bombay, with minors in Computer Science and Machine Intelligence, where he focused on autonomous driving with Amit Sethi.

- **Lovish Madaan (PhD student at UCL/Meta)**: Lovish is a PhD researcher in Artificial Intelligence at Meta AI and UCL NLP, focusing on thinking models, synthetic data, and improved evaluation methods. Previously, he worked at Google Research India in the Machine Learning and Optimization team with Prateek Jain and Srinadh Bhojanapalli, where he explored inference-efficient machine learning and natural language processing. He holds a B.Tech and M.Tech in Computer Science and Engineering from the Indian Institute of Technology, Delhi, where he worked with Parag Singla, Sayan Ranu, and Aaditeshwar Seth on a range of research projects.

- **Kurt Keutzer (Professor at UC Berkeley)**: Kurt is a Professor Emeritus in the Department of Electrical Engineering and Computer Science at the University of California, Berkeley, where he has been a faculty member since 1998. He received his Ph.D. in Computer Science from Indiana University in 1984. Prior to his academic career, he worked at AT&T Bell Laboratories and served as Chief Technical Officer and Senior Vice President of Research at Synopsys, Inc.. At Berkeley, his research focuses on systems issues related to the application of deep learning in areas such as computer vision, speech recognition, natural language processing, and finance. He has co-authored six books and over 250 refereed articles, and was elected a Fellow of the IEEE in 1996. In 2015, he co-founded DeepScale, a company specializing in efficient deep neural networks for automotive applications, which was acquired by Tesla in 2019.

- **Prateek Jain (Director, Google DeepMind India)**: Prateek is a Principal Scientist and Director at Google DeepMind India, where he leads the Machine Learning and Optimization team. He earned his Ph.D. from UT Austin and his B.Tech from IIT Kanpur. His research focuses on machine learning, efficient and elastic inference of large models, retrieval and indexing, rigorous non-convex optimization and reinforcement learning. Prateek regularly serves on the senior program committees of top machine learning conferences and is on the editorial boards of leading journals including JMLR and SIMODS. He has received multiple best paper awards, including the 2020 Best Paper Award from the IEEE Signal Processing Society, as well as the Young Alumnus Award from IIT Kanpur in 2021 and the ACM India Early Career Researcher Award in 2022.

5.1 Advisors

We are grateful to have Ion Stoica (Professor at UC Berkeley), Inderjit Dhillon (VP, Google Fellow at Google and Professor at UT Austin) and Mihir Sanjay Kale (Research Engineer at Meta) as advisors for our workshop.

6 Previous Related Workshops and Novelty

Our proposed **Workshop on Scaling Post-training for LLMs (SPOT)** builds upon a series of recent ICLR workshops that have explored aspects of large-scale learning, reasoning, and alignment in foundation models. While these prior venues focused on applications or reliability, none systematically addressed the **scaling science of post-training itself** — the study of how compute, data, and algorithms interact to determine post-training efficiency and asymptotic performance.
Among the most relevant prior workshops are:

- Reasoning and Planning for Large Language Models — examined reasoning and decision-making abilities of post-trained models, motivating the need to understand how such capabilities emerge with scale.
- Reliable and Responsible Foundation Models — focused on the trustworthiness and robustness of foundation models, emphasizing post-training as a key stage for alignment.
- Data-centric Machine Learning Research (DMLR) — explored how data scaling laws and data curation impact performance, a perspective crucial for understanding data efficiency in post-training.
- Open Science for Foundation Models — promoted transparency, reproducibility, and open evaluation in foundation model research, aligning closely with our workshop's commitment to shared post-training benchmarks and reproducible scaling studies.

- Building Trust in LLMs and LLM Applications — investigated evaluation, interpretability, and safety practices for aligned LLMs, which intersect with the verification and stability dimensions of scalable post-training.

These workshops collectively highlight the growing recognition that post-training is central to the reliability and performance of modern LLMs. However, their primary emphasis has been on *applications* and *evaluation*, rather than the underlying scaling dynamics of post-training. The **SPOT** workshop aims to bridge this gap by bringing together diverse perspectives from academic and industrial researchers and practitioners, to share practical experiences, and to outline a clear research direction toward building a principled science of post-training at scale.

7 Access, Advertising and Outreach

To ensure broad visibility and active participation, we will implement a coordinated outreach strategy combining web presence, visual media, and social engagement. We have a dedicated workshop website (https://spoticlr.github.io/) which will serve as the central hub for all information, including the call for papers, submission instructions, speaker details, and schedule, and will later host recordings and accepted papers to promote long-term accessibility. We will design a professional digital poster and social media banners to circulate across academic mailing lists and research groups (e.g., BAIR, FAIR, DeepMind, AI2, etc). The call for submissions and speaker announcements will be widely shared on X (Twitter) and LinkedIn through the organizers' and invited speakers' accounts, collectively reaching a broad audience of researchers and practitioners in large-scale learning and reinforcement learning. Through these combined efforts, we expect to attract a diverse and engaged audience spanning academia, industry, and open research communities.

References

Milad Aghajohari, Kamran Chitsaz, Amirhossein Kazemnejad, Sarath Chandar, Alessandro Sordoni, Aaron Courville, and Siva Reddy. The markovian thinker, 2025. URL https://arxiv.org/abs/2510.06557.

Zhipeng Chen, Xiaobo Qin, Youbin Wu, Yue Ling, Qinghao Ye, Wayne Xin Zhao, and Guang Shi. Pass@k training for adaptively balancing exploration and exploitation of large reasoning models, 2025. URL https://arxiv.org/abs/2508.10751.

Devvrit, Sneha Kudugunta, Aditya Kusupati, Tim Dettmers, Kaifeng Chen, Inderjit Dhillon, Yulia Tsvetkov, Hannaneh Hajishirzi, Sham Kakade, Ali Farhadi, and Prateek Jain. Matformer: Nested transformer for elastic inference, 2024. URL https://arxiv.org/abs/2310.07707.

Hongzhe Du, Weikai Li, Min Cai, Karim Saraipour, Zimin Zhang, Himabindu Lakkaraju, Yizhou Sun, and Shichang Zhang. How post-training reshapes llms: A mechanistic view on knowledge, truthfulness, refusal, and confidence, 2025. URL https://arxiv.org/abs/2504.02904.

[@_katieeverett] Katie Everett. Power Laws Tweet. X, May 2025. URL https://x.com/_katieeverett/status/1925665335727808651.

William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity, 2022. URL https://arxiv.org/abs/2101.03961.

Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun Mei, Jiashu Wang, Tongkai Yang, Binhang Yuan, and Yi Wu. Areal: A large-scale asynchronous reinforcement learning system for language reasoning, 2025. URL https://arxiv.org/abs/2505.24298.

Tyler Griggs, Sumanth Hegde, Eric Tang, Shu Liu, Shiyi Cao, Dacheng Li, Charlie Ruan, Philipp Moritz, Kourosh Hakhamaneshi, Richard Liaw, Akshay Malik, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. Evolving skyrl into a highly-modular rl framework, 2025. Notion Blog.

Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. OLMo: Accelerating the science of language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.841. URL https://aclanthology.org/2024.acl-long.841/.

Horace He and Thinking Machines Lab. Defeating nondeterminism in llm inference. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20250910. https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/.

Xinyu He and Xianghui Cao. Understanding post-training structural changes in large language models, 2025. URL https://arxiv.org/abs/2509.17866.

Prime Intellect. Intellect 2 release. https://www.primeintellect.ai/blog/intellect-2-release, 2025. Accessed: 2025-10-10.

Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnawamy, and Rohan Kadekodi. Diskann: Fast accurate billion-point nearest neighbor search on a single node. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/09853c7fb1d3f8ee67a61b6bf4a7f8e6-Paper.pdf.

Margaret Li, Sneha Kudugunta, and Luke Zettlemoyer. (mis)fitting: A survey of scaling laws, 2025. URL https://arxiv.org/abs/2502.18969.

Meta AI. The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation, 2025. URL https://ai.meta.com/blog/llama-4-multimodal-intelligence/.

Mika, Sami, Johannes, and Prime Intellect. Planetary-scale inference: Previewing our peer-to-peer decentralized inference stack. https://www.primeintellect.ai/blog/inference, 2025.

Johan Obando-Ceron, Ghada Sokar, Timon Willi, Clare Lyle, Jesse Farebrother, Jakob Foerster, Gintare Karolina Dziugaite, Doina Precup, and Pablo Samuel Castro. Mixtures of experts unlock parameter scaling for deep rl, 2024. URL https://arxiv.org/abs/2402.08609.

OpenAI. Introducing OpenAI o3 and o4-mini. https://openai.com/index/introducing-o3-and-o4-mini/, 2025. Accessed: 22 September 2025.

Alex Piche, Rafael Pardinas, Ehsan Kamalloo, and Dzmitry Bahdanau. Pipelinerl. 2025. URL https://huggingface.co/blog/ServiceNow/pipelinerl.

John Schulman and Thinking Machines Lab. Lora without regret. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20250929. https://thinkingmachines.ai/blog/lora/.

Zelin Tan, Hejia Geng, Mulei Zhang, Xiaohang Yu, Guancheng Wan, Yifan Zhou, Qiang He, Xiangyuan Xue, Heng Zhou, Yutao Fan, Zhongzhi Li, Zaibin Zhang, Guibin Zhang, Chen Zhang, Zhenfei Yin, and Lei Bai. Scaling behaviors of llm reinforcement learning post-training: An empirical study in mathematical reasoning, 2025. URL https://arxiv.org/abs/2509.25300.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kenealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Petrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szpektor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Põder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. Gemma 3 technical report, 2025. URL https://arxiv.org/abs/2503.19786.

xAI Team. Grok 4. 2025. URL https://x.ai/news/grok-4.

Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction, 2025. URL https://arxiv.org/abs/2408.15240.

Haizhong Zheng, Jiawei Zhao, and Bedi Chen. Prosperity before collapse: How far can off-policy rl reach with stale data on llms?, 2025. URL https://arxiv.org/abs/2510.01161.

Xuekai Zhu, Daixuan Cheng, Dinghuai Zhang, Hengli Li, Kaiyan Zhang, Che Jiang, Youbang Sun, Ermo Hua, Yuxin Zuo, Xingtai Lv, Qizheng Zhang, Lin Chen, Fanghao Shao, Bo Xue, Yunchong Song, Zhenjie Yang, Ganqu Cui, Ning Ding, Jianfeng Gao, Xiaodong Liu, Bowen Zhou, Hongyuan Mei, and Zhouhan Lin. Flowrl: Matching reward distributions for llm reasoning, 2025. URL https://arxiv.org/abs/2509.15207.