Exploration in AI Today: Sponsors Overview

The Exploration in AI Today Workshop at ICML 2025 (exait-workshop.github.io) tackles one of the most pressing challenges in modern AI: how to efficiently collect observations for optimization, control, and generalization in vast, complex search spaces. As traditional exploration struggles to scale, this workshop reimagines its role in the era of foundation models—highlighting where smart exploration can outperform brute-force data and compute, as seen in breakthroughs like DeepSeek-R1. With a focus on real-world impact across robotics, LLM alignment, and scientific discovery, the workshop brings together diverse voices to explore how prior knowledge, uncertainty quantification, and modern learning techniques (like diffusion and in-context learning) can unlock the next wave of AI progress. Held at ICML—one of the premier conferences in machine learning—the workshop will attract top researchers from both academia and industry. We anticipate over 200 attendees, drawn by the timely topic and a standout lineup of speakers. The workshop is scheduled to set place on July 19th in Vancouver, Canada, and is organized as a collaboration between ETH Zurich, UC Berkeley and Google DeepMind (Appendix B).

Commitment to Diversity, Equity, and Inclusion. Our organizing team and invited speakers reflect our efforts to include women and members of underrepresented groups, and the same standard will be applied to panelists. Further, we are offering travel grants and registration fee waivers to make the workshop accessible to a larger community. We also expect a diverse audience across technical disciplines and backgrounds, driven by the workshop's interdisciplinary scope.

Use of Funds. Sponsor contributions will directly support the following elements.

- Paper Awards: The workshop accepts papers that focus on the topic of exploration in four tracks of theory and RL, Language Modeling, AI for Science, and Robotics. In each track, we intend to award $1000 to the best paper. This amounts to a total of $4000 in expenses.
- Catering: We are planning to provide a coffee break with refreshments during the day and an evening apéro to encourage informal discussion and networking. We estimate the costs at $2,500.
- Recognition dinner: Following customs for workshops in ML, we will invite all speakers and organizers for a dinner after the workshop to acknowledge and thank them for their contributions. We expect 15 participants to attend this dinner and estimate the costs at $1,500.
- Attendees Travel Support: We are aiming to cover travel costs for distinguished speakers and offer a small travel and registration support grant to students from underrepresented communities. We expect funds to cost up to $10,000.

Sponsorship Tiers. Depending on the interest of the sponsors in supporting specific elements of the event (e.g., paper awards, apéro, or DEI travel support), we can create a custom package. Follows our general proposition for sponsorship tiers.

Tier	Contribution	Benefits
Gold	$10,000	Logo on all workshop materials and website, acknowledgment in opening remarks, reserved entry passes for affiliates, talk and poster slots, visibility as the key supporter of the awards.
Silver	$5,000	Logo on website, mention in opening remarks, 1 free pass for company affiliates.
Bronze	$2,000	Logo on website.

Sponsors. We are supported by the Swiss National Center of Competence in Automation Research (NCCR) and ETH AI Center contributing $5,000 and $2,500 at Silver and Bronze tiers, respectively.

A Additional Information on the Workshop

How can we efficiently collect observations for optimization, control, and generalization? This is a key challenge in AI and is known as the exploration problem. Effective exploration has driven progress in areas such as robotics, recommender systems, and scheduled medical trials. However, as we address larger, more complex applications—such as drug discovery or language modeling—the exceptionally large search spaces render traditional exploration algorithms ineffective. As a result, recent breakthroughs in AI have come not from traditional exploration algorithms, but largely from training large foundation models on diverse corpora of pre-existing, curated datasets. Despite this, we have witnessed sparks showing that exploration, when done right, can compensate for data and computation—for example, in the training of DeepSeek-R1—suggesting that exploration can still play a key role in AI today.

The Exploration in AI Today (EXAIT) Workshop at ICML 2025 will focus on addressing the evolving role of exploration in AI. We will dwell on the question: what is the place of exploration in today's AI landscape and in which settings can exploration algorithms address current open challenges? In particular, we consider the potentially pivotal role that exploration might play in navigating complex and high-dimensional search spaces across real-world applications such as robotics, large language model alignment, and AI for science. We seek to answer fundamental questions such as:

1. How might we leverage prior knowledge—for example, in the form of a pretrained model or physics simulator—to enable efficient exploration, and how might biases present in such prior knowledge affect the resulting exploration?
2. How can we effectively explore high-dimensional search spaces? In particular, how might we exploit the structure inherent in many high-dimensional settings of interest to enable effective exploration?
3. Uncertainty quantification historically has played a critical role in exploration: how can we quantify uncertainty effectively to direct exploration in high-dimensional settings such as language?
4. How can modern learning techniques (e.g. diffusion, in-context learning) be utilized in exploration?

We aim to leverage insights from such questions to address application-driven challenges, for example:

1. How can we develop agents—LLM agents, robots, and beyond—that effectively explore their environment online to determine how to accomplish a desired goal?
2. Which experiments should we run in order to collect the most informative data to update our models—whether a language model, world model, model of biological interactions, or more?
3. What role can exploration play in reasoning, e.g., in searching and generating chains of thought?

We will seek diverse perspectives from across areas to provide insights into these questions and allow for cross-pollination between application areas and theoretical advances—enabling researchers to understand how successful exploration techniques in other fields can address open problems in their own. We believe that this will spur both the development of fundamental exploration techniques, as well as the advancement of AI across domains, from robotics, to LLMs, the natural sciences, and beyond.

Organization and Schedule. The workshop will offer invited and contributed talks, a panel discussion, a poster session, a small best paper award ceremony, and a sponsored apéro, to foster different levels of interactions between the attendees. The timetable gives an overview of the day, and the following paragraph explains the idea behind every event.

We have organized **the invited talks** into three themed sessions. A long morning session on advances of exploration in theory, RL, and Robotics, followed by two application-focused sessions on exploration for language modeling and AI for science. These sessions align with the paper submission areas and **the best paper awards**. We invite researchers to submit their recent works demonstrating the role of exploration in four categories: 1) theory and exploration methodology, 2) robotics, 3) LLM post-training, and 4) AI for science (including but not limited to drug discovery, materials design, physical sciences, etc). For each category we award a paper with a **contributed talk**, certificate, and prize money. **The poster session**, which effectively runs throughout the day and has a 2 hour dedicated slot, will allow for extended discussions on

all accepted work. The **panel discussion** will be an encounter of academics and industry researchers, whose expertise covers a spectrum from exploration at scale to theoretical foundations of exploration. We will curate the panelists line-up at a later date, partly from the invited speakers and partly from among prominent community members who will decide to attend ICML in person.

Morning sessions	8:50-9	Opening remarks
	9-10:30	Invited talks by Alison, Sergey & Dorsa (Theory, RL, and Robotics session)
	10:30-11:30	Panel discussion
	11:30-14	Posters and lunch
Afternoon sessions	14-15:30	Invited talks by Ben, Dylan and Rishabh (Language modeling session)
	15:30-16	Coffee break
	16-17	Invited talks by Ji Won and Masatoshi (AI for science session)
	17-18	Announcement of the best papers and contributed talks (4 papers, one talk each)
	18-19	Closing remarks and sponsored apéro

List of Speakers. The following speakers are confirmed to attend in person:

Sergey Levine (UC Berkeley/Physical Intelligence)
Alison Gopnik (UC Berkeley)
Benjamin Van Roy (Stanford)
Dorsa Sadigh (Stanford)
Rishabh Agarwal (McGill/Google DeepMind)
Dylan Foster (Microsoft Research)
Masatoshi Uehara (Wisconsin-Madison)
Ji Won Park (Genentech).

Speaker selection. We carefully curated a list of speakers to ensure a broad representation of diverse perspectives and applications of exploration algorithms in AI. All our invited speakers are leading experts in exploration methods such as reinforcement learning (RL) and focus on diverse applications like robotics, language models, and drug discovery. Specifically, Benjamin Van Roy, Rishabh Agarwal, and Dylan Foster focus on leveraging RL algorithms for fine-tuning large language models (LLMs). Moreover, Benjamin and Dylan are also known for their contributions to fundamental algorithmic theory. Dorsa Sadigh and Sergey Levine are recognized for their research on exploration and robotics. Masatoshi Uehara and Ji Won Park are renowned for their work on AI-driven drug design and lab-in-the-loop data collection strategies. Lastly, Alison Gopnik brings a unique perspective, drawing from her extensive research on how infants learn and explore, and comparing those strategies to the learning approaches used in modern AI systems.

In addition to representing a variety of research backgrounds, our speakers and organizers reflect a range of regional backgrounds, research areas, career stages, and perspectives from both academia and industry. Several of the workshop organizers and speakers are members of groups that are underrepresented in ML and STEM, including women, and people of color researchers.

Previous and Related events. This will be the **first edition of the EXAIT workshop**. Venues like ICML and NeurIPS have hosted workshops with various themes related to exploration, including: (*i*) exploration in RL (ICML 18, 19), (*ii*) language and RL (NeurIPS 22), (*iii*) foundation models for decision-making (NeurIPS 23), (*iv*) adaptive experimental design and active learning in the real world (ICML 22, NeurIPS 23), (*v*) foundations of RL and control (ICML 24), (*vi*) aligning RL experimentalists and theorists (ICML 24), and (*vii*) Bayesian decision-making and uncertainty (NeurIPS 24). These workshops either target specific applications (e.g., language, foundation models) or provide a platform for dialogue between different subfields, such as active learning, RL, or uncertainty quantification.

On the contrary, the EXAIT workshop focuses on the crucial and timely challenge of exploration in complex and high-dimensional yet structured spaces, and how exploration methods can be scaled, adapted and integrated in today's AI landscape. The workshop will bring together researchers from diverse fields, including algorithms research, robotics, large language models (LLMs), and drug discovery, to share their insights and innovative approaches. It provides a unique opportunity for cross-disciplinary collaboration, fostering in-depth discussions on exploration algorithms and their applications.

B List of Organizers and Advisors

This workshop is being managed by a motivated team of 6 organizers and 4 advisors who have collectively organized over 20 workshops at top machine learning venues such as ICML, NeurIPS, ICLR, and CORL. The advisors are: Olivier Bachem (Director, Research Scientist at Google DeepMind), Pieter Abbeel (Professor at UC Berkeley), Andreas Krause (Professor at ETH Zurich), and Kevin Jamieson (Associate Professor at University of Washington). The junior organizers are:

Parnian Kassraie (🌐, ✉, 𝔊) Parnian is a PhD student at ETH Zurich, advised by Prof. Andreas Krause and a visiting scholar at Carnegie Mellon University. Her research develops theoretical foundations for structure-aware online learning and optimization, with applications in LLM fine-tuning and drug discovery. She is a recipient of the 2024 Google PhD Fellowship.

Dr. Andrew Wagenmaker (🌐, ✉, 𝔊) Andrew is a postdoctoral researcher at UC Berkeley and Berkeley Artificial Intelligence Research (BAIR) working with Prof. Sergey Levine. He obtained his PhD at the University of Washington, where he was advised by Prof. Kevin Jamieson. His research centers on developing learning-based algorithms for decision-making in sequential environments, both in theory and practice, and has spanned settings such as multi-armed bandits, reinforcement learning, and continuous control.

Bhavya Sukhija (🌐, ✉, 𝔊) Bhavya is a PhD at ETH Zurich, advised by Prof. Andreas Krause and Prof. Stelian Coros. His research lies at the intersection of reinforcement learning (RL) and robotics. He has worked on the development, theoretical analysis, and real-world deployment of RL algorithms across several different settings, including online RL, active/intrinsic exploration, non-episodic RL, safe exploration, and meta RL. His work has been featured in top venues for ML and robotics.

Dr. Carmelo (Carlo) Sferrazza (🌐, ✉, 𝔊) Carlo is postdoctoral researcher at UC Berkeley and Berkeley Artificial Intelligence Research (BAIR), where he is advised by Prof. Pieter Abbeel. His research focuses on advancing humanoid intelligence by incorporating priors, inductive biases, and multi-sensory feedback. Carlo received his PhD at ETH Zurich where he was advised by Prof. Raffaello D'Andrea. During his Ph.D. research, he worked on the design of vision-based, data-driven tactile sensors, and the applications of such sensors to robot control and dexterous manipulation. His work has been featured in top ML and robotics venues and TEDX talks.

Lenart Treven (🌐, ✉, 𝔊) Lenart is a PhD at ETH Zurich where he is advised by Prof. Florian Dörfler and Prof. Andreas Krause. The core focus of his PhD, is continuous-time, time-adaptive, and efficient reinforcement learning.

Amy X Lu (🌐, ✉, 𝔊) Amy is a PhD student at UC Berkeley and Berkeley Artificial Intelligence Research (BAIR), where she is advised by Prof. Pieter Abbeel. She is also a part-time researcher on the machine learning team at Genentech. The core focus of Amy's research is AI for drug discovery, particularly via multimodal generation and biological foundation models.