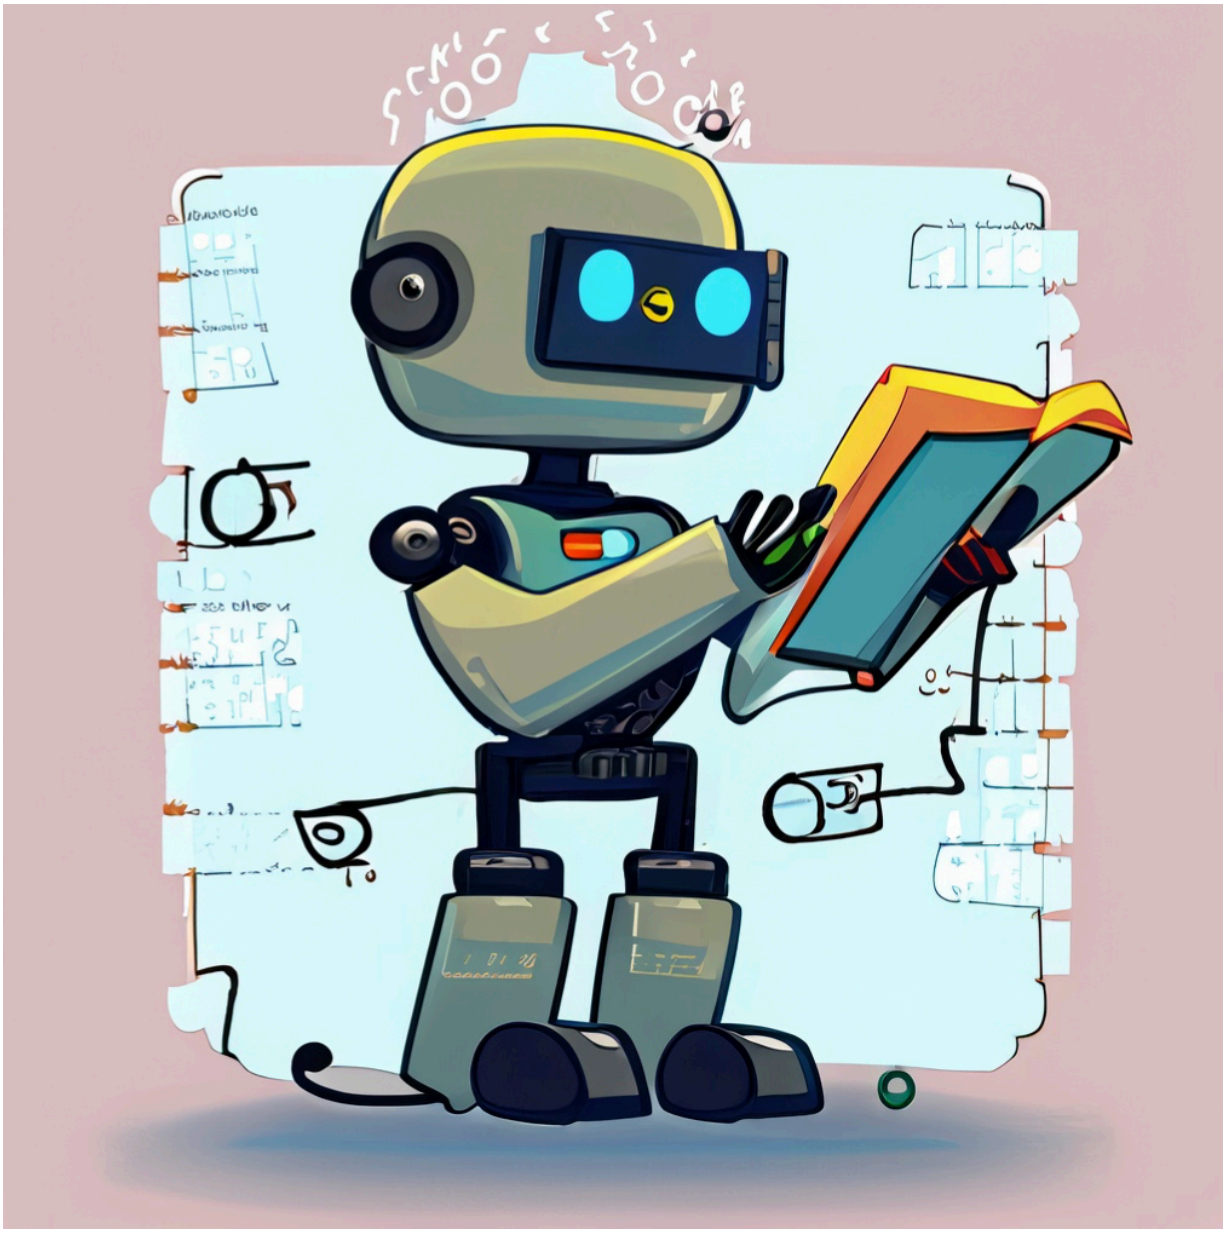

Generative Models for Decision Making

ICLR'24 Workshop Proposal

Abstract

Generative Artificial Intelligence (AI) has made significant advancements in recent years, particularly with the development of large language and diffusion models. These generative models have demonstrated impressive capabilities in various tasks, such as text generation [1-3] and image and audio synthesis [4, 5, 6,7]. Concurrently, Reinforcement Learning (RL) has made significant strides in solving complex sequential decision-making problems with the help of external knowledge sources [8-10]. However, there remains untapped potential in combining generative models with RL algorithms to tackle real-world challenges, particularly to improve sample efficiency of *tabula rasa* training by introducing priors from related domains such as visual question-answering, image captioning and image generation.

This workshop aims to bring together researchers and practitioners from the fields of generative AI and reinforcement learning to explore the latest advances, methodologies, and applications. By fostering collaborations between these two domains, we intend to unlock new opportunities for addressing complex problems that lie at the intersection of both fields.

The workshop will cover a wide range of topics, including but not limited to:

- **Large Language Models and Reinforcement Learning:** Exploring how large language models, such as GPT-3 and beyond, can be integrated with RL algorithms to improve performance on complex sequential decision-making tasks. Moreover, we welcome contributions which study how to make large language model suitable for interactive and embodied settings, be it for planning, reward generation, simulation of the physical world or introducing human priors into RL via language. Tentative research questions: which benchmarks, evaluation criteria and environments should be developed by the community to assess the utility of large language models for RL?
- **Diffusion Models and Reinforcement Learning:** Investigating the potential of diffusion models and other generative models for enhancing RL algorithms for

image generation, reinforcement learning from pixels, and robotic control. Tentative research questions: can diffusion models be used as physics-aware world models, thus improving the sample efficiency of online RL methods?

- **Sample Efficiency in Reinforcement Learning:** Discussing techniques for improving sample efficiency in RL through generative models, enabling the application of RL in data-constrained environments. Specifically, can generative models trade reward-labelled efficiency by using more unlabelled samples? Tentative research questions: can we use large language model or video prediction models to enable faster learning on complex, open-ended RL tasks?
- **Exploration in Reinforcement Learning:** Exploring how generative models can facilitate exploration strategies in RL, especially in high-dimensional and sparse reward settings. For instance, since generative models can efficiently represent parts of the data distribution, it is reasonable to assume that they can also provide an informative learning signal. Tentative research questions: how can pre-trained generative models help RL agents solve long-horizon, sparse reward or open-ended tasks without a clear definition of success?
- **Transfer Learning in Reinforcement Learning with Generative Models:** Investigating methods to leverage pre-trained generative models for transfer learning in RL, enabling agents to adapt to new tasks more efficiently through a deeper understanding of the underlying dynamical system of RL problems. Tentative research questions: do generative models used for high-level planning or low-level control transfer better to unseen domains than classical RL methods?
- **Inverse Reinforcement Learning and Imitation Learning:** Analyzing how generative models can assist IRL/IL algorithms in learning from observed behaviour, or used for data augmentation. Tentative research questions: can generative models capture richer information contained in human demonstrations than existing methods?

Format:

The workshop will consist of a mix of invited talks, contributed paper presentations, panel discussions, and interactive sessions. Researchers are encouraged to submit

original research papers, work-in-progress reports, and position papers that highlight the synergies between generative AI and RL.

Intended Audience:

The workshop targets AI researchers, industry professionals, and graduate students interested in generative models and RL. Participants should have a basic understanding of generative models and reinforcement learning concepts. Familiarity with recent advancements in both fields will be beneficial but not mandatory.

Invited Speakers

Noam Brown (OpenAI) - *Confirmed*

Noam Brown is a Research Scientist at OpenAI working on multi-step reasoning, self-play, and multi-agent AI. Previously, he worked at FAIR (Meta), where he developed CICERO, the first AI to achieve human-level performance in the strategy game Diplomacy.

Igor Mordatch (Google DeepMind) - *Confirmed*

Igor Mordatch is a Research Scientist at Google DeepMind working on foundation models for decision making. Throughout his career, he has worked on emergent abilities in multi-agent systems; decision making as conditional sequence modeling; and optimal control and reinforcement learning techniques for robotics.

Katja Hofmann (Microsoft Research) - *Confirmed*

Katja Hofmann is a Senior Principal Researcher at Microsoft Research Cambridge. She leads a team that focuses on Deep Reinforcement Learning for Games, with a mission to advance the state of the art in reinforcement learning, driven by current and future applications in video games.

Joelle Pineau (FAIR, McGill)

Joelle Pineau is an Associate Professor and William Dawson Scholar at McGill University where she co-directs the Reasoning and Learning Lab. Member of Mila's faculty corp, she also leads the Facebook AI Research lab in Montreal, Canada. She works on applying planning and learning algorithms to complex problems in robotics, health care, games and conversational agents.

Jeanette Bohg (Stanford) - *Confirmed*

Jeanette Bohg is a Professor for Robotics at Stanford University and directs the Interactive Perception and Robot Learning Lab. Her research explores two questions in robotic grasping and manipulation: What are the underlying principles of robust sensorimotor coordination in humans, and how can we implement them on robots?

Karthik Narasimhan (Princeton) - *Confirmed*

Karthik Narasimhan is an Assistant Professor of Computer Science at Princeton University and co-director of Princeton Natural Language Processing. His research spans the areas of natural language processing and reinforcement learning.

Schedule

08:45	Introduction and opening remarks
09:00	Invited talk 1
09:30	Invited talk 2
10:00	Contributed Talk (Oral)
10:30	**Coffee Break + Posters**
11:30	Contributed Talk (Oral)
12:00	Invited talk 3
12:30	**Lunch Break**
14:00	Invited talk 4
14:30	Invited talk 5
15:00	Contributed talks (3 spotlights)
15:30	**Coffee break + Posters**
16:30	Invited talk 6
17:00	**Panel discussion**
18:00	**Socials**

Submission guidelines

Researchers are invited to submit extended abstracts or full papers in PDF format. Submissions should follow the main track ICLR format and be limited to the same page length as the main track. We encourage the use of LaTeX templates for submissions. Detailed submission guidelines and templates will be provided on the workshop website.

Publication:

Accepted papers will be made available on the workshop website. The workshop will not have any proceedings.

Organizer Bios

Bogdan Mazoure is a Research Scientist in the Machine Learning Research (MLR) team at Apple. He obtained his Ph.D. in Computer Science at McGill University and Quebec AI Institute (Mila) in 2022, during which he spent time at Microsoft Research, Google Brain and Google DeepMind, working on reinforcement learning and generative models. He previously organized the ICLR 2021 workshop on Self-Supervised Learning for RL, and the ICML 2021 workshop on Theory of Continual Learning.

Devon Hjelm is a Research Scientist with MLR at Apple and an External Industry Affiliate at Mila. He earned his PhD focusing on generative models for Neuroimaging data at the University of New Mexico, then did a PostDoc under Yoshua Bengio at Mila studying representation learning and mutual information estimators. He authored and co-authored several seminal works on self-supervised learning, contrastive learning, and mutual information estimation as a Principal Researcher at Microsoft Research (MSR). He currently focused on representation learning in interactive and embodied settings as well as using LLMs in Reinforcement Learning.

Lisa Lee is a Research Scientist at Google DeepMind, and currently works on Gemini. She obtained her Ph.D. in Machine Learning at Carnegie Mellon University on deep reinforcement learning. Lisa has served as Workflow Chair for ICML 2019, and has organized workshops on Ecological Theory of Reinforcement Learning (NeurIPS 2021); Learning with Rich Experience: Integration of Learning Paradigms (NeurIPS 2019); and Theoretical Foundations and Applications of Deep Generative Models (ICML 2018).

Roberta Raileanu is a Research Scientist at Meta and an Honorary Lecturer at UCL. She earned her PhD in Computer Science from NYU where she worked on generalization in deep reinforcement learning. Currently, she works on augmenting foundation models with planning and decision making abilities by training them from feedback and interaction with external tools, environments, and humans. Roberta has previously co-organized the Workshop on Unsupervised RL (URL) at ICML 2021, the Agent Learning in Open-Endedness (ALOE) at ICLR 2022 and NeurIPS 2023, as well as the workshop on Socially Responsible Language Models Research (SoLaR) at NeurIPS 2023.

Yilun Du is a PhD student at MIT. Currently, he works on constructing intelligent robotic decision making systems through the use of generative models. He has previously co-organized the Foundation Models for Decision Making workshop (FMDM) at NeurIPS 2022 and 2023.

Walter Talbott is a Research Scientist at Apple. He received his Ph.D. from UC San Diego, focused on visual perception and control of a pneumatic humanoid robot. He is currently focused on deep reinforcement learning, robotics, and generative models.

Alexander Toshev is a Research Scientist and Manager at Apple ML Research. He holds a Doctorate in Computer and Information Sciences from the University of Pennsylvania, and Diplom-Inform. from University of Karlsruhe, Germany. His work lies in Computer Vision, Robotics, Multimodal Learning, and Embodied AI. More recently he co-initiated the use of foundational models in Robotics applications. He is a co-organizer on multiple workshops, CVPR Workshop on Visual Semantic Navigation, a series of CVPR Workshops on Embodied AI.

Katherine Metcalf is a Research Scientist at Apple. She received her Ph.D. from Indiana University focusing on adapting methods of self-supervised learning in humans to machine learning systems. She is currently focused on using foundation models to improve agent abilities to learn from humans.

Funding

Upon acceptance, we will solicit sponsorships from companies such as Apple, Meta and Google to improve the accessibility and inclusivity of our event. In particular, we will cover the registration fees of participants who indicate they won't be able to attend otherwise, prioritizing members from under-represented groups.

Diversity Statement

During the selection of organizers and speakers, we actively encouraged all forms of diversity. The expertise of invited speakers ranges from game theory to natural language processing and robotics, while also achieving gender parity and covering different research institutions, nationalities, scientific backgrounds and levels of seniority. The full scale of scientific seniority is covered in both organization committee and speakers, including PhD candidates as well as senior researchers.

Similar Past Events

Generative models such as large language models and diffusion-based models have received a lot of attention lately, with workshops focusing on the broad topic, algorithmic innovations or a specific sub-domain. In particular, *Efficient Systems for Foundation Models* (ICLR 2023) accepted submissions studying efficiency of foundation models. The theoretical properties of foundation models were also studied in submissions to the *Mathematical and Empirical Understanding of Foundation Models* (ICLR 2023).

Other workshops, such as *Generative AI and Law* (ICLR 2023) and *Foundation Models for Decision Making* (FMDM, NeurIPS 2022), studied the application of generative models to a specific domain. While the FMDM workshop focuses on how we can build foundation models for decision making tasks, our workshop focuses more how we can use existing generative AI models and techniques such as language or diffusion models to improve the sample efficiency, exploration, transfer, imitation, of RL agents. In addition, we plan to discuss how RL can improve generative AI models and what new developments are needed to do so.

References

[1] Hoffmann, Jordan, et al. "Training compute-optimal large language models." 2022.

[2] Alayrac, Jean-Baptiste, et al. "Flamingo: a visual language model for few-shot learning." 2022.

[3] OpenAI. "GPT-4V(ision) system card." 2023.

[4] Rombach, Robin, et al. "High-resolution image synthesis with latent diffusion models." 2022.

[5] Ho, Jonathan, et al. "Imagen video: High definition video generation with diffusion models." 2022.

[6] Oord, Aaron van den, et al. "Wavenet: A generative model for raw audio." 2016.

[7] Liu, Haohe, et al. "Audioldm: Text-to-audio generation with latent diffusion models." 2023.

[8] Wang, Guanzhi, et al. "Voyager: An open-ended embodied agent with large language models." 2023.

[9] Lifshitz, Shalev, et al. "STEVE-1: A Generative Model for Text-to-Behavior in Minecraft." 2023.
[10] Baker, Bowen, et al. "Video pretraining (vpt): Learning to act by watching unlabeled online videos." 2022.