

Sharath Chandra Raparthy

London, UK
sharathraparthy@gmail.com

sharathraparthy.github.io

Areas of interest: Large Language Models, Reinforcement Learning.

EXPERIENCE

Reka AI September 2023 — present
Member of Technical Staff London

- Drove core development of Reka Research, an AI agent that answers complex questions by dynamically browsing the web and private documents.
- Led Supervised Fine-Tuning (SFT) effort for the Reka Research agent, establishing the training foundation for subsequent reinforcement-learning fine-tuning.
- Delivered state-of-the-art results on SimpleQA (91.7% accuracy) and ResearchEval (59.1% accuracy) beating competitors like Perplexity, OpenAI and Anthropic.
- Advanced Reinforcement-Learning fine-tuning for Reka Flash 3, our open-source reasoning model that gets best-in-class performance compared to competitors.

Meta - Fundamental AI Research (FAIR) August 2022 — August 2024
AI Resident London

- **Tool-use and Math Reasoning for Llama 3**
 - Core contributor to tool-use fine-tuning efforts for Llama-3 models.
 - Responsible for integration of tools like web search, code execution, Wolfram|Alpha and File Uploads which improved the reasoning and data analysis capabilities.
 - Designed and implemented supervised fine-tuning and rejection-sampling pipelines, improving Llama=3's mathematical reasoning accuracy.
- **Safety**
 - Co-lead Rainbow Teaming, a method that detects vulnerabilities in LLMs with a 92% attack success rate while maintaining high quality and diversity of generated adversarial prompts.
 - Pioneered the integration of Rainbow Teaming into Meta's automatic red-teaming pipelines and also Llama 3 safety finetuning.
 - Presented this work at Neural Information Processing Systems (NeurIPS 2024)
- **In-context Learning for Sequential Decision Making**
 - Spearheaded a transformer-based in-context learning approach enabling zero-shot adaptation from just a handful of demonstrations on MiniHack and Procgen tasks.
 - Demonstrated substantial gains on out-of-distribution benchmarks through curriculum design and model-data scaling.
 - Presented at the International Conference for Machine Learning (ICML) 2024 and featured on the TalkRL podcast.
- **LLM Reasoning and Refinements**
 - Developed GLoRe, a novel approach that proposes global and local refinement approach to improve LLM reasoning.
 - GLoRe improved the performance on GSM8K of a strong RL finetuned LLama-2 13B model by 12%, demonstrating the effectiveness of the proposed global and local refinement techniques.

Recursion Pharmaceuticals March 2022 — Aug 2022
Machine Learning Research Intern Montreal

- Worked on a project to generate diverse candidates for Multi-Objective Optimization, crucial for drug discovery and material design.
- **Contributed to the design and implementation of Multi-Objective GFlowNets (MOGFNs)**, a novel method for diverse candidate generation, resulting in an **ICML publication**.
- Conducted extensive experiments on synthetic and benchmark tasks, demonstrating the effectiveness of our proposed methods, particularly in improving candidate diversity.

Montreal Institute for Learning Algorithms (Mila) July 2019 — August 2020
Research Intern | Advisors: Liam Paull and Irina Rish Montreal

- Addressed the 'sim-2-real' transfer challenge in robot learning, devising two unique methods grounded in curriculum learning.
- Formulated an algorithm targeting the resolution of the negative adaptation issue prevalent in meta-reinforcement learning.
- Carried out comprehensive experiments on actual robots, resulting in three workshop papers presented at **ICLR and RSS**.

EDUCATION

Master of Science (MSc) in Artificial Intelligence,
Mila/University of Montreal
Advisor: Prof. Irina Rish
Thesis: On Impact of Mixing Times in Continual Reinforcement Learning
Nominated for Dean's List

Sept 2020 — Aug 2022

Bachelor of Technology (B.Tech) in Mechanical Engineering,
National Institute of Technology (NIT), Nagpur

Sept 2015 — May 2019

PUBLICATIONS

- [1] AI@Meta. “Llama 3 Model Card”. In: (2024). URL: https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- [2] Abhimanyu Dubey et al. *The Llama 3 Herd of Models*. 2024. arXiv: 2407.21783 [cs . AI]. URL: <https://arxiv.org/abs/2407.21783>.
- [3] Alex Havrilla, Yuqing Du, **Sharath Chandra Raparthi**, Christoforos Nalmpantis, Jane Dwivedi-Yu, Maksym Zhuravinskyi, Eric Hambro, Sainbayar Sukhbaatar, and Roberta Raileanu. *Teaching Large Language Models to Reason with Reinforcement Learning*. 2024. arXiv: 2403.04642 [cs . LG].
- [4] Alexander Havrilla, **Sharath Chandra Raparthi**, Christoforos Nalmpantis, Jane Dwivedi-Yu, Maksym Zhuravinskyi, Eric Hambro, and Roberta Raileanu. “GLoRe: When, Where, and How to Improve LLM Reasoning via Global and Local Refinements”. In: *Forty-first International Conference on Machine Learning*. 2024. URL: <https://openreview.net/forum?id=LH6R06NxdB>.
- [5] **Sharath Chandra Raparthi**, Eric Hambro, Robert Kirk, Mikael Henaff, and Roberta Raileanu. “Generalization to New Sequential Decision Making Tasks with In-Context Learning”. In: *Forty-first International Conference on Machine Learning*. 2024. URL: <https://openreview.net/forum?id=IVQ4FUZ6dp>.
- [6] Mikayel Samvelyan*, **Sharath Chandra Raparthi***, Andrei Lupu*, Eric Hambro, Aram H. Markosyan, Manish Bhatt, Yuning Mao, Minqi Jiang, Jack Parker-Holder, Jakob Foerster, Tim Rocktäschel, and Roberta Raileanu. “Rainbow Teaming: Open-Ended Generation of Diverse Adversarial Prompts”. In: *Advances in Neural Information Processing Systems*. Vol. 37. 2024. URL: https://proceedings.neurips.cc/paper_files/paper/2024/file/8147a43d030b43a01020774ae1d3e3bb-Paper-Conference.pdf.
- [7] Moksh Jain, **Sharath Chandra Raparthi**, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Yoshua Bengio, Santiago Miret, and Emmanuel Bengio. “Multi-Objective GFlowNets”. In: *ICML (2023)*. arXiv: 2210.12765 [cs . LG].
- [8] Sarthak Mittal, **Sharath Chandra Raparthi**, Irina Rish, Yoshua Bengio, and Guillaume Lajoie. “Compositional Attention - Disentangling Search and Retrieval”. In: *ICLR Spotlight (Top 5%) (2022)*.
- [9] Matthew Reimer*, **Sharath Chandra Raparthi***, Ignacio Cases, Maxmillian Puelma Touzel, and Irina Rish. “Continual Learning in Environments with Polynomial Mixing Times”. In: *CoRR abs/2112.07066 (2021)*. arXiv: 2112.07066 [cs . LG].
- [10] Bhairav Mehta, Tristan Deleu*, **Sharath Chandra Raparthi***, Christopher Pal, and Liam Paull. “Curriculum and Benchmarking in Meta-Reinforcement Learning”. In: *ICLR BeTR-RL workshop (2020)*.
- [11] **Sharath Chandra Raparthi**, Bhairav Mehta, Florian Golemo, and Liam Paull. “Generating Automatic Curricula via Self-Supervised Active Domain Randomization”. In: *ICLR BeTR-RL workshop (2020)*.
- [12] **Sharath Chandra Raparthi**, Melissa Mozifian, Liam Paull, and Florian Golemo. “CuNAS - CUriosity-driven Neural-Augmented Simulator”. In: *Sim2Real Transfer for Robotics Workshop at RSS (2020)*.
- [13] Akshay Kulkarni, Navid Panchi, **Sharath Chandra Raparthi**, and Shital Chiddarwar. “Data Efficient Stagewise Knowledge Distillation”. In: *Arxiv (2019)*.
- [14] Anil Kota, **Sharath Chandra Raparthi**, Parag Khanna, and Torjborn Dahl. “Explicit Sequence Proximity Models for Hidden State Identification”. In: *Reinforcement Learning Under Partial Observability Workshop at NeurIPS (2018)*.

TECHNICAL SKILLS

Programming	Python
Libraries	PyTorch, NumPy
Others	Github, Markdown, LaTeX, Shell Scripting, SLURM.

COURSES

Probabilistic Graphical Models (A+), Reinforcement Learning and Optimal Control (A+), Game theory and Machine Learning (A+) and Dynamical Systems (A), Data Structures (A), Introduction to Algorithms (A+), Numerical Linear Algebra (A+)

AWARDS AND VOLUNTEERING/MENTORSHIP

- 2022** Serving as a Program Chair for [Machine Learning Reproducibility Challenge \(MLRC\) 2022](#)
- 2021** Served as an Area Chair for [Machine Learning Reproducibility Challenge \(MLRC\) 2021](#)
- 2021** Received Thesis Excellence Scholarship of \$2500 from University of Montreal to support my masters research
- 2021** Served as a reviewer for NeurIPS' 21, Preregistration Workshop and Ecological Theory in RL at NeurIPS
- 2020** Volunteered for Mila buddy program for Fall 2020 and 2021 cohort and helped on-boarding 4 students to the Mila ecosystem
- 2020** Mentored 5 students in deep learning projects at [lvLabs](#), the largest student-run robotics lab of India
- 2019** Volunteered [Reinforcement Learning and Decision Making](#) conference (RLDM - 2019) in Montreal