

PROFESSIONAL SUMMARY

RL researcher focused on deploying foundation models in real world systems. Published at top-tier venues (NeurIPS, ICML), experienced in post-training VLAs, profiling-first performance engineering, and GPU optimization across the training + inference stack. Demonstrated expertise in translating research advances into production systems, and accelerating inference for real time deployment constraints.

EDUCATION

MILA QUEBEC

Research MSc in Computer Science (GPA: 4.15/4.30)

Montreal

Sept 2021 - Nov 2023

- **Thesis:** Sample efficient reinforcement learning for real world systems. Supervised by Samira E. Kahou

PROFESSIONAL EXPERIENCE

ARAYA INC

Deep RL Researcher

Tokyo

Jan 2024 - Present

- Building **post training and deployment frameworks for VLAs** for human robot interactions in assistive robotics tasks.
- Demonstrated that simple offline actor-critic RL can post-train a billion-scale VLA, **improving task success by 29 pp** without specialized flow objectives, distillation, or test-time guidance.
- Built an optimized batched inference engine for VLAs that enabled online RL by **reducing latency by 100x (20s to 0.2s)**.
- Optimized internal pipelines (**2x faster** training and inference), enabling same day data collection-training-deployment.
- Built automated data annotation pipelines for training **embodied reasoning agents** from internal datasets.

MICROSOFT RESEARCH

Research Intern. Supervisor: John Langford

New York

June 2023 - Sept. 2023

- Developed a novel reachability-aware latent representation using multi-step inverse dynamics for sample-efficient RL.
- Enabled computationally efficient hierarchical planning, achieving **3x faster** goal reaching and **50% fewer operations**.
- Demonstrated significant improvements in sample efficiency for reward-based settings and enabled zero-sample ad hoc goal reaching in reward-free environments.

RESEARCH EXPERIENCE

IMPROVING SAMPLE EFFICIENCY THROUGH THE REDUCIBLE LOSS: MILA, QUEBEC

Supervisor: Samira E. Kahou

Montreal

May. 2022 - Jan 2023

- Devised the Reducible Loss (ReLo) prioritization that improves performance by **20%** over vanilla prioritized experience replay (PER), increasing the sample efficiency and accuracy of the value function. Accepted to **NeurIPS 2023**.
- ReLo can be integrated in **any** off policy algorithm with minimal code changes and negligible computational overhead.

REPRESENTATION LEARNING FOR MODEL BASED RL: MILA, QUEBEC

Supervisor: Samira E. Kahou, Danijar Hafner (Google DeepMind)

Montreal

Jan. 2021 - Feb 2022

- Improved sample efficiency of training with **2x** fewer gradient steps through Variational Sparse Gating (VSG) by inducing the ability to sparsely update changes in the GRU. Accepted at **NeurIPS 2022**.
- Built the BringBackShapes environment that tests agents on long horizon tasks with partial observability and stochasticity.

PUBLICATIONS

Resolving Uncertainty in Robotic Manipulation with Vision-Language Models and Conformal Prediction

Preprint

Shogo Akiyama*, **Shivakanth Sujit***, Luca Nunziante*, Rousslan Fernand Julien Dossa, and Kai Arulkumaran

Open-loop VLM Robot Planning: An Investigation of Fine-tuning and Prompt Engineering Strategies

Preprint

Shogo Akiyama, Rousslan Fernand Julien Dossa, Kai Arulkumaran, **Shivakanth Sujit**, Edward Johns

PcLast : Discovering Plannable Continuous Latent States: Paper

ICML 2024

Anurag Koul*, **Shivakanth Sujit***, Shaoru Chen, and Ben Evans, Lili Wu, Riashat Islam, Raihan Seraj, Yonathan Efroni, Miroslav Dudik, John Langford, Alex Lamb

Prioritizing Samples in Reinforcement Learning with Reducible Loss: Paper Code

NeurIPS 2023

Shivakanth Sujit, Somjit Nath, Pedro Braga, Samira E Kahou