

Mohamed Ayman Mohamed

Continual reinforcement learning. Research scientist and software engineer.

✉ mamoham3@ualberta.ca

🎓 Google Scholar

🔗 github.com/mayman150

🌐 LinkedIn

HIGHLIGHTS

- **Published** at **RLC**, **ICLR**, and **TPAMI**, with a fourth paper under review. **63** citations, h-index **3** ([Scholar](#)). MSc in Computing Science (AI), University of Alberta, **4.0/4.0**, supervised by Marlos C. Machado.
- **Research that ships**: first author of **AgarCL**, a continual RL environment and benchmark released as a maintained C++/Python library, now used to study loss of plasticity.
- **Production engineering** at AWS scale, currently applying generative models to Kubernetes and region infrastructure operations.
- **Unusual range**: comfortable writing CUDA kernels, profiling C++ with Valgrind, and designing distributed services, having also built a pipelined RISC-V core and an x86-64 bootloader from scratch.

RESEARCH EXPERIENCE

Graduate Research Assistant, Reinforcement Learning

Apr 2024 – Aug 2025

University of Alberta, Machado Lab *Full-time*

- Built a C++/Python continual RL environment modelling **gradual** and **abrupt** distribution shift in episodic and non-episodic settings, across **15** simulation scenarios.
- Benchmarked **DQN**, **PPO**, and **SAC** on policy stability, convergence speed, and sample efficiency, exposing a **90%** failure rate under continual conditions.
- Roughly doubled throughput with CUDA rendering, asynchronous inference, and quad-tree approximation: grid observations from **1,200** to **2,200** steps/second, screen from **800** to **1,600**.
- Built a CNN feature-extraction pipeline learning compact state representations from high-dimensional inputs: **30%** faster convergence, **25%** better sample efficiency.
- Containerised with Docker behind a CI/CD pipeline of **25+** automated tests; resolved memory leaks with Valgrind and tracked regressions with Google Benchmark.

Machine Learning Research Intern

Feb 2023 – Oct 2023

KAUST, King Abdullah University of Science and Technology *Remote*

- Implemented a chain-of-thought ground-truth generation tool, improving 3D object detection accuracy by **8%** on complex visual grounding tasks.
- Built a rule-based anchors parser for automated object localisation, published at **ICLR 2024**, with an NLP pipeline (tokenisation, TF-IDF, semantic filtering) that lifted anchor to target matching precision **15%**.
- Optimised **3DComPaT++** segmentation with **PointNet++** and diffusion-based refinement, deployed a full-server solution, and improved IoU by **10%**.

Undergraduate Research Assistant, Machine Learning

Sep 2022 – Apr 2023

American University in Cairo *Part-time*

- Designed an inverted-index query execution engine for disk-based approximate nearest-neighbour search, reaching a **6×** speedup over Microsoft's SPANN at comparable recall and graph quality.
- Improved retrieval accuracy **15%** with autoencoder and transformer embeddings, and cut end-to-end latency **30%** on billion-scale datasets via automated embedding and index-construction pipelines.

WORK EXPERIENCE

Software Engineer

Sep 2025 – Present

Amazon Web Services *Full-time*

- Reduced AWS region build-model creation from **days to about five minutes** with an AI-powered code generation system on Amazon Bedrock.
- Initiated an AI automation program for EKS operational workflows, targeting the highest-volume repetitive work: build-model authoring, pipeline failure triage, ticket resolution, and dashboard analysis.
- Architected cross-partition observability infrastructure and adaptive data pipelines, cutting operational overhead **40%** and enabling zero-touch deployments.
- Delivered automated testing frameworks, compliance workflows, and a resilient resource cleanup API, reducing customer escalations **25%**.

Software Engineer

Oct 2024 – Apr 2025

University of Alberta *Contract*

- Built a distributed simulation module in C# and C++ against the High-Level Architecture standard, cutting communication latency by **20%**.

- Designed a PostgreSQL schema managing **500k+** records across **6** federates, behind an N-tier ASP.NET service handling **500+** concurrent requests at sub-second latency.
- Embedded the CPython interpreter in C# so Python analysis and ML workflows could run inside .NET.

Applied Scientist Intern

Jul 2022 – Sep 2022

Microsoft, Zcode Team *Full-time*

- Developed a meta-learning data distillation algorithm cutting dataset size by **99.9%** while holding **95%** of model performance across transformer architectures.
- Benchmarked across NLP datasets, improving **10%** over the state of the art in machine translation, classification, and named-entity recognition.
- Showed across **10** controlled experiments that model non-linearity significantly affects performance on distilled data.

Software Engineer

Sep 2021 – Apr 2022

Suitera L.L.C. *Full-time*

- Found the bottleneck in slab data processing for complex circuit and multi-die package simulations, then wrote C++ approximation algorithms that simplify the geometry while preserving the structure that matters.
- Built aggregation for similar slab segments to cut simulation runtime, validated with custom GoogleTest fixtures, profiled with gdb and Valgrind, and wrote OpenGL 3D visualisations of circuit structure and current flow.

Full-Stack Software Engineering Intern

Jun 2021 – Aug 2021

Envelope.network Inc. *Remote*

- Architected a cross-browser extension raising engagement **15%**, plus a **5**-stage recommendation pipeline for **500+** users at **0.76** precision, served by **10+** Scala/Finatra endpoints over PostgreSQL.

PUBLICATIONS

Self-Referenced Social Preferences: Cooperation without Observing Others' Rewards 2026

Mohamed A. Mohamed, Harshil Kotamreddy, Marcos Menon José

Under review. Preprint: [arXiv:2610.07881](https://arxiv.org/abs/2610.07881)

The Cell Must Go On: Agar.io for Continual Reinforcement Learning 2026

Mohamed A. Mohamed, Kateryna Nekhomiazh, Vedant Vyas, Marcos M. José, Andrew Patterson, Marlos C. Machado

Reinforcement Learning Conference (RLC). [arXiv:2505.18347](https://arxiv.org/abs/2505.18347). Code: [machado-research/AgarCL](https://github.com/machado-research/AgarCL)

3DCoMPaT++: An Improved Large-Scale 3D Vision Dataset for Compositional Recognition 2025

Habib Slim, Xiang Li, Yuchen Li, Mahmoud Ahmed, **Mohamed Ayman**, et al.

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)

CoT3DRef: Chain-of-Thoughts Data-Efficient 3D Visual Grounding 2024

Eslam Abdelrahman, **Mohamed Ayman**, Mahmoud Ahmed, Habib Slim, Mohamed Elhoseiny

International Conference on Learning Representations (ICLR)

EDUCATION

MSc, Computing Science (Artificial Intelligence) Sep 2023 – Aug 2025

University of Alberta

GPA 4.0/4.0. Supervisor: Prof. Marlos C. Machado. Travel Award 2024/25, Graduate Research Assistantship.

BSc, Computer Engineering, Minor in Mathematics Sep 2018 – May 2023

American University in Cairo

GPA 3.97/4.0. Dean's List four consecutive years. Al-Ghurair Scholarship, 2018 to 2023. Bronze Medal, ECPC 2023; ACPC Finalist 2023.

TEACHING

Graduate TA, Alberta Fundamentals of Reinforcement Learning, Introduction to Machine Learning, Fundamentals of Databases, Introduction to Networks

Undergraduate TA, AUC Introduction to Machine Learning, Deep Learning in Practice, Analysis & Design of Algorithms, Fundamentals of Computing II

TECHNICAL SKILLS

Languages	Python, C/C++, C#, Java, Scala, Rust, JavaScript, SQL
Machine learning	PyTorch, JAX, TensorFlow, PyTorch Lightning, Hugging Face, scikit-learn, XGBoost, Transformers, spaCy, OpenCV, PointNet++
Reinforcement learning	CleanRL, PFRL, Stable-Baselines3, RLlib, Gymnasium
Cloud & infrastructure	AWS (Bedrock, CDK, Lambda, EC2, S3), Azure, Docker, Kubernetes, CI/CD, PostgreSQL, gRPC
Low-level & tooling	CUDA, OpenMP, OpenGL, x86-64 assembly, RISC-V, Verilog, CMake, Valgrind