

ASIM OSMAN

📍 Cape Town, South Africa ✉️ asim@aims.ac.za 🌐 github.com/Asimawad

PROFILE

I am a Google DeepMind Scholar and a Research Engineer with a Masters degree in AI from the University of Cape Town (UCT). I have experience in both research and industry. My primary interests lie in Deep Learning, Large Language Models (LLMs), and Reinforcement Learning. At present, I am focused on optimizing the performance of reasoning LLMs and advancing the capabilities of LLM agents in coding and machine learning engineering tasks by scaling inference-time compute and inference-time scaling strategies.

EDUCATION

MSc. in Artificial Intelligence for Science (AI for Science) 2024 – 2025

University of Cape Town – African Institute for Mathematical Sciences (AIMS), Cape Town, South Africa

Graduated with Distinction — **Top Student (AIMS Excellence Award)**

Curriculum containing a heavy mathematical background coupled with exciting courses on the use of AI in various scientific disciplines.

Thesis: *Scaling Inference-Time Compute for Machine Learning Engineering Agents*

Supervisor: **Arnu Pretorius** (Stellenbosch University)

[Thesis PDF](#) [Certificate](#)

BSc. (Honours) in Electrical and Electronic Engineering / Software Engineering 2016 – 2023

University of Khartoum, Sudan

Final GPA: 9.18 / 10.0 (Best Graduation Project Award - Joint First)

Thesis Topic: Benchmarking Neural Machine Translation for Arabic-Swahili

Released the largest dataset for Machine Translation between Arabic and Swahili. Proposed using a small dataset to fine-tune a large model for translating between Arabic and Swahili.

Published at ICLR 2023, Tiny Papers Initiative — *Machine Translation Baselines for Arabic-Swahili*

Sudanese Secondary School Certificate (SSSC) 2013 – 2016

Alitisam Model Secondary School, Kassala, Sudan

Score: 96.1% (7th nationwide)

EXPERIENCE

Research Engineer Intern Jan 2025 – Jan 2026

InstaDeep, Cape Town, South Africa

- Developed autonomous research agents using large language models to automate machine learning workflows.
- Applied advanced inference-time scaling strategies to improve large-scale multi-agent reinforcement learning systems.

Research Engineer Intern / Research Assistant May 2024 – Aug 2024

Enigma AI Research and Development Team, Khartoum, Sudan

- Worked on the development of agriAI, a field monitoring application providing real-time insights into crop health, soil conditions, and environmental factors.

Data Science Intern Jul 2022 – Mar 2023

AmunData, Data-Driven Research, and Advisory (Technical Team), Khartoum, Sudan

- Built a crop classification pipeline using Random Forest with Sentinel-2 multispectral bands, achieving 89.22% accuracy in categorizing key crops in El Gezira.
- Contributed to the agriAI project — an intelligent farming assistant providing real-time crop health and soil analysis.

PROJECTS

AIDE: AI-Driven Exploration for ML Engineering — GitHub — Thesis

An open-source, autonomous agent that tackles end-to-end ML engineering tasks by framing development as a *solution space tree search*. Powered by LLMs, AIDE drafts initial solutions, debugs faulty code, and iteratively improves working scripts—mirroring the workflow of a human data scientist.

Key elements and results:

- Implements inference-time scaling (ITS) strategies—*self-consistency*, *self-reflection*, and *decomposed planner-coder (task decomposition)*—to enhance smaller, open-source models.
- Core loop: *Solution Generator* (LLM), *Evaluator* (executes code and parses outputs/metrics), and a heuristic *Search Policy* to select the next node for refinement.
- Evaluated distilled DeepSeek-R1 models (7B/14B/32B) with a high-throughput vLLM backend on a curated subset of MLE-Bench competitions.
- The best agent—**32B + Decomposed Planner-Coder**—achieved a **30%** medal rate (Pass@6), *matching* OpenAI o4-mini and *surpassing* the GPT-4-Turbo baseline on our subset.

Fine-Grained Credit Assignment for RLHF (VinePPO) — GitHub

Designed and implemented **Fine-Grained Credit PPO**, a JAX-based reinforcement learning framework that assigns token-level rewards to stabilize LLM alignment training.

Key contributions: migration from PyTorch to JAX/Flax, TPU-optimized training, modular API-based credit assignment using GPT-4/Gemini judges.

Arabic-Swahili Neural Machine Translation — GitHub

Created the first large-scale Arabic-Swahili translation dataset and fine-tuned multilingual transformers for low-resource translation tasks (ICLR 2023). Developed and fine-tuned models for Arabic to Swahili translation, overcoming data scarcity challenges and establishing benchmarks for future research in these languages.

Active-Learning-Core-Concepts

Demonstrated the principles of active learning using the Forest Covertypes Dataset and various query strategies.

Comparing-Uncertainty-Methods-Active-Learning

Explored uncertainty quantification in deep learning models using information theory.

Bayesian-Deep-Active-Learning

Conducted experiments with Bayesian neural networks for active learning.

MovieLens Large Scale Recommender System

Implemented a movie recommendation system using matrix factorization techniques based on user ratings, leveraging latent factor models to capture patterns in user behavior and movie characteristics. Built a scalable recommendation system leveraging collaborative filtering.

Vehicle-Re-Identification

Designed a system using ResNet50 and triplet loss for vehicle re-identification with the VeRi dataset.

Fully-Convolutional-Networks-for-Image-Denoising

Implemented fully convolutional networks to denoise images from the LFWcrop dataset.

Computer-Vision-Transfer-Learning

Investigated transfer learning for CIFAR-10 classification, enhancing understanding with saliency maps.

Reinforcement-Learning-LunarLander-REINFORCE

Applied the REINFORCE algorithm to solve the LunarLander-v2 environment using policy gradients.

AWARDS AND SCHOLARSHIPS

- **2025:** AIMS Excellence Award for the best academic performance in the MSc in AI for Science program at AIMS, South Africa.
- **2024:** Google DeepMind Scholarship to study MSc in AI for Science at AIMS, South Africa.
- **2023:** Best Graduation Project (Joint First). This is a yearly prize awarded by the University of Khartoum to the creators of the best graduation project in the Engineering Faculty.
- **2022:** Best Student Award, Department of Electrical and Electronic Engineering, University of Khartoum (GPA: 9.54).
- **2016:** Government bursary for the first 100 highest achievers to finance undergraduate degree at the University of Khartoum.
- **2013:** State scholarship for high achievers to finance high school education, Kassala, Sudan.

PUBLICATIONS

Papers and Preprints:

- Asim Osman, Ahmed Elmahdy, Hiba Imam, ICLR Conference 2023, Tiny Papers Initiative Machine Translation Baselines for Arabic-Swahili, 2023

TALKS, CONFERENCES, WORKSHOPS, AND RESEARCH ACTIVITIES

- **2025:** Deep Learning Indaba AI Conference, Rwanda (Attendee & Organizer).
- **2024:** AI/ML Paper Reading Club at AIMS South Africa (Member).
- **2023:** Deep Learning Indaba AI Conference, Ghana (Inaugural Reviewer & Short Paper Submission Reviewer).
- **2022:** IndabaX Sudan (Content Team Leader, Organizer, Speaker). Gave the Talk: Introduction to PyTorch.
- **2017 & 2018:** Electrical and Electronic Engineering Student Exhibition (Presented a Project).

TECHNICAL SKILLS

Languages & Frameworks: Python (Advanced), PyTorch (Advanced), JAX/Flax (Beginner/Intermediate), Optax, NumPy (Advanced), Pandas (Advanced), Scikit-learn (Intermediate), Seaborn/Matplotlib (Intermediate)

Specialties: LLMs, RLHF (PPO, GRPO, DPO), Distributed Training, vLLM Inference, Inference-Time Scaling, TPU Optimization, Huggingface (Advanced)

Tools: Git/GitHub (Advanced), Linux, Bash, Docker (Intermediate), Cloud Computing, Power BI/Excel (Intermediate)

LANGUAGES

Arabic: Native

English: Professional Working Proficiency (B2/C1)