HIRAD EMAMIALAGHA

GRONINGEN, NETHERLANDS +31 6 1559 28 35

DETAILS

Groningen
Netherlands
+31 6 1559 28 35
hirad.alagha@outlook.com

LINKS

My LinkedIn Page

SKILLS

- Artificial Intelligence and Robotics
- Machine Learning Algorithms
- Software Engineering and Architecture Design
- MLOps & DevOps
- Infrastructure Engineering
- Integration Engineering
- Amazon AWS
- Databricks
- Apache Spark, Polars, Pandas
- Git / CI-CD
- Docker, ECS, Kubernetes, Nomad
- Python
- MLFlow
- EDD, BDD, TDD
- Scrum, SAFe, Agile
- Deep Learning
- Generative Models: Diffusion, VAE, GAN, LLM
- Kafka, Benthos, MSK
- DataDog, Qlik, Sofa, Ganglia, Grafana, Cloudwatch, Prometheus
- Django, Flask, FastAPI

PROFILE

Senior MLOps & AI Engineer with a strong foundation in artificial intelligence, robotics, and large-scale ML systems. Experienced in designing and operating production grade ML platforms, while maintaining a deep interest in reinforcement learning, generative models, and autonomous systems. Comfortable owning complex platforms end-to-end, leading technical direction, and raising engineering standards across teams.

Personal Traits

Fast learner, Committed, Adaptive, Critical Thinker, Problem Solver, Organized

EMPLOYMENT HISTORY

Senior Lead MLOps Engineer at PlayerUnknown Production, Amsterdam
Oct 2025 — Present

- Designing and leading a next-generation ML research and production platform supporting large-scale experimentation, GPU workloads, and reproducible research
- Owning platform architecture decisions across orchestration, storage, observability, and ML lifecycle management
- Enabling ML researchers to move from ad-hoc experimentation to structured, production-ready workflows
- Acting as technical lead at the intersection of ML research, infrastructure, and engine teams
- Architected and implemented the underlying ML infrastructure from scratch using Nomad + Consul + Temporal, enabling GPU-aware scheduling across heterogeneous nodes, service discovery, and reliable execution of long-running ML workloads
- Drove cross-team planning with ML research, engine, and operation teams, owning workload prioritization, GPU capacity planning, hardware requirements, and continuous optimization of platform utilization, reliability, and operational cost

MLOps Engineer at Essent, 's-Hertogenbosch
Oct 2022 — Oct 2025

- Took the lead in rework of our most important predictive model, directly having a positive impact our entire customer base.
- Managed the end-to-end development life-cycle, leading a team of engineers and fostering collaboration across multiple departments.
- Achieved significant optimization in performance against the legacy pipelines, reducing daily runtime by over 85%.
- Expanded cross-team test coverage and refined end-to-end integration and deployment strategies, enhancing overall project reliability.
- Boosted model accuracy and consistency by elevating feature quality and systematically reducing technical debt.
- Developed a sophisticated, in-depth monitoring and operational framework that significantly boosted performance consistency, strengthened platform resilience, and increased stakeholder confidence.

LANGUAGES

English

Persian

Dutch

MLOps Engineer - Scrum Master at Essent, 's-Hertogenbosch
Dec 2021 — Oct 2022

- Led Essent's main data science team in their transformation journey to become an end-to-end team capable of independently delivering high quality models to production.
- Led Essent in building MLOps capabilities by introducing MLOps way of working principle, designing & developing the underlying frameworks, and building the expertise in the team by mentoring the members.
- Implemented a framework around Databricks, MLFlow, Delta Lake and DataDog for productionalization and monitoring of our models
- Took the responsibility of the team's scrum master to establish the team's way of working according to the company's established SAFe framework.
- Responsible in the development and maintenance of various models, and corresponding feature pipelines ranging from predicting customer energy usage to indicating solar panel installations.

Application Development Analyst at Accenture, Utrecht
Mar 2020 — Dec 2021

- Implemented, productionalized, and maintained various data and model pipelines using different AWS services such as SageMaker, EC2, Glue, Batch, Lambdas, Apache Airflow, and Apache NiFi.
- Significantly improved the performance of our pipelines by reworking legacy implementations in Scala, RapidMiner, and SageMaker notebooks to optimized pipelines using Python, Spark, Docker, and other tools.
- Developed a custom Data Quality Monitoring system (DQM) for monitoring various data and feature qualities across dozens of pipelines.
- Designed and implemented GDPR solutions for big data on AWS S3 data lake, covering hundreds of tables and thousands of partitions.
- Developed and automated many unit, integration, and stress tests on various GitLab CI/CD pipelines..
- Actively participated in operational tasks, maintaining and updating a number of data and model pipelines on a daily basis.
- Won the award for Accenture's International compition yearly "Do-Good" competition for Accenture Netherlands - 2020.

EDUCATION

Master's Degree in Artificial Intelligence and Robotics, University of Groningen (RUG), Groningen Netherlands
Feb 2017 — Feb 2019

Courses: Machine Learning, Handwriting Recognition, Multi-Agent Systems, Robotics, Signals and Systems, Cognitive Engineering, Cognitive Robotics, Computer Vision, Natural language processing, and Modeling Complex Behavior.

Bachelor's Degree in Computer Science, Sunway University, Subang Jaya-Malaysia
Aug 2012 — Jun 2016

Dual Programme with Lancaster Unviersity - United Kingdom

Courses: Object-Oriented Programming, Embedded Systems, Data Structures & Algorithms, Human-Computer Interaction, Artificial Intelligence, Computational Intelligence, Software Architecture, and Design Patterns, Distributed Systems, Computer Networks, Operating System Fundamentals.

PUBLICATION

Communicating Intention in Decentralized Multi-Agent Multi-Objective Reinforcement Learning Systems, Groningen Netherlands
Jul 2019

Supervised by prof. Marco Wiering, this project extended the state-of-the-art in decentralized multi-agent reinforcement learning by introducing "*Intention*" mechanism using multi-objective reinforcement learning and enabled independent Q-learners to communicate their intentions in order to enhance the coordination in their individual behaviors. The RL agents were trained and tested on Groningen's Peregrine HPC cluster and were 42 parallel computation clusters ran more than 10 million simulations. The newly introduced method almost doubles the convergence rate from the baseline approaches and significantly reduces the learning time of the agents as the complexity of the task increased.

Paper found at: https://api.semanticscholar.org/CorpusID:235669388

PromptCraft: The Ultimate Guide to AI Image Generation, Amsterdam
Aug 2025

Published a technical and creative guide book focused on AI image generation, detailing prompt engineering strategies, generative model characteristics, and practical workflows for controlling and optimizing visual outputs from diffusion based models.

Amazon Kindle Link: https://a.co/d/bpIWPfS

COURSES & CERTIFICATES

Microsoft Azure Certified: Azure Data Engineer Associate , Microsoft
May 2020 — Apr 2023

Microsoft DP-200: Implementing an Azure Data Solution, Microsoft
May 2020 — Apr 2023

Microsoft DP-201: Designing an Azure Data Solution , Microsoft
May 2020 — Apr 2023

Machine Learning in Production Course, Databricks
Jan 2022 — Mar 2022

HOBBIES

Thanks to my background and hobby as an indie game developer, I spend a significant amount of my free time learning about and creating music, developing video games, designing graphic models, creating websites and mobile apps, and exploring new robotic and AI projects.

In addition, I have a passion for traveling abroad, biking, playing musical instruments, and watching documentaries.

REFERENCES

References available upon request