

ARAM GHOLIKIMILAN

London, UK | hello@itsaram.dev | itsaram.dev |
linkedin.com/in/aram-gholikimilan | github.com/itsaram-ai

PROFESSIONAL SUMMARY

AI/ML engineer with an MSc in Artificial Intelligence from King's College London and a First Class Honours BSc in Computer Science from City, University of London, with two departmental prizes. I build LLM agent systems with retrieval grounding, structured-output validation and automated evaluation, and run controlled machine learning experiments end to end, most recently benchmarking a 12-qubit quantum autoencoder across 800,000+ ICU waveform windows.

TECHNICAL SKILLS

Languages: Python, Java, C++, C#, SQL, Dart

AI and LLM systems: LLM APIs (Anthropic), agent orchestration, retrieval-augmented generation, TF-IDF retrieval, structured output validation (pydantic), LLM evaluation harnesses, golden test sets, groundedness checking, prompt versioning, decision tracing

Machine learning: PyTorch, TensorFlow, Keras, deep learning, CNNs, LSTMs, temporal convolutional networks, autoencoders, reinforcement learning (PPO, Stable-Baselines3, Q-learning), Markov decision processes and value iteration, SVM, decision trees, clustering (k-means, hierarchical), Naive Bayes, NLP text preprocessing and stemming, hyperparameter optimisation, computer vision

Research and evaluation: Controlled benchmarking, ROC-AUC and F1 evaluation, K-fold cross-validation, leakage and generalisation testing, out-of-distribution evaluation, quantum machine learning (variational circuits, amplitude encoding, swap test, Classiq)

Data and infrastructure: pandas, NumPy, scikit-learn, SciPy, matplotlib, NLTK, Firebase (Firestore, Auth, Functions), Google Cloud Run, Google Cloud Services, Azure VM, Git and GitHub, pytest, Wireshark

EXPERIENCE

Creative-Director | Multi-Agent LLM Content Pipeline

Jun 2026 | Open-source project | Python, Anthropic API, pydantic, pytest

- Built a three-stage planner, writer and critic pipeline that generates, critiques and revises content against a retrieval-grounded knowledge base.
- Implemented retrieval-augmented generation with a from-scratch TF-IDF retriever and chunk-level citation of every claim.
- Enforced structured JSON output with pydantic and a validation-retry loop, scored drafts on a four-dimension critic rubric, and versioned every prompt.
- Built an automated evaluation harness with a golden test set covering schema validity, platform constraints and citation accuracy, 18/18 checks passing.

Quantum Autoencoders for Sensor Diagnostics | MSc Dissertation

Feb 2026 - Jul 2026 | King's College London, in collaboration with Classiq | Python, PyTorch, Classiq

- Built and ran the circuits on the Classiq quantum platform, with the encoding scheme shaped by technical discussions with a Classiq engineer.
- Benchmarked a 12-qubit variational quantum autoencoder against LSTM and temporal-convolutional baselines for anomaly detection on ICU waveforms.
- Found the quantum model competitive with the convolutional baseline (ROC-AUC 0.6682 against 0.6701) using 54 trainable parameters against 36,038.
- Built the full pipeline in Python: windowing, quality filtering, patient-level splits and five synthetic anomaly types across 800,000+ waveform windows.
- Verified the local simulator against the Classiq platform to one part in 10^{15} , and reproduced results under 4,096-shot sampling to within 0.0010.

MNIST Convolutional Network and Robustness Evaluation Harness

Feb 2026 | King's College London | Python, TensorFlow, Keras, NumPy

- Designed and trained a 1.2M-parameter convolutional network in TensorFlow and Keras, reaching 99.67% accuracy on the MNIST test set.
- Built the network with batch normalisation, in-graph augmentation, global average pooling and AdamW, with learning-rate scheduling and early stopping.
- Wrote a 554-line evaluation harness to test the model beyond the brief, covering four out-of-distribution datasets and misclassified-sample export.
- Verified against leakage with a shuffled-label test at chance accuracy, and measured a 0.11% train-holdout to test generalisation gap.

Fitness and Nutrition Mobile App | BSc Individual Project

Jan 2025 - Apr 2025 | Awarded the Outstanding BSc Project Prize | Flutter, Dart, Firebase, PPO, Google Cloud Run

- Built a cross-platform fitness and nutrition app in Flutter and Firebase with auth, tracking and nutrition scoring.
- Integrated a PPO model generating personalised training and nutrition plans, deployed on Google Cloud Run.
- Cut app load times through persistent model loading, API caching and Firestore indexing.

Machine Learning Pipeline for Protein Localisation

Nov 2024 - Dec 2024 | Python, scikit-learn

- Built and evaluated models for protein cellular localisation on a balanced yeast dataset using K-fold cross-validation.
- Tuned an SVM with RandomizedSearchCV across three kernels, reaching 66.54% test accuracy against a 25% baseline.

Markov Decision Process Solver for Pacman

Oct 2025 - Dec 2025 | King's College London | Python, standard library only

- Modelled the game as a Markov decision process, solved by value iteration on every move under a stochastic motion model.
- Met a no-external-libraries constraint with a local-window state space, BFS reward shaping and a per-move time budget.
- Scored 35/35 and 20/20 on the two hardest automated layout tests, and 93/100 across all four.

Face Covering Detection

Jan 2025 - Apr 2025 | Python, OpenCV, scikit-learn

- Compared three models (HOG+SVM, HOG+MLP and a custom CNN), reaching 91% test accuracy on unseen real-world images.
- Added face detection, histogram equalisation and augmentation, and shipped a MaskDetection inference function.

Data Mining: Classification, Clustering and Text Mining

Feb 2026 - Mar 2026 | King's College London | Python, pandas, scikit-learn, NLTK

- Built decision-tree income classification on the UCI Adult dataset with missing-value handling and one-hot encoding.
- Compared k-means, k-means++ and agglomerative clustering over three cluster counts and ten restarts by silhouette score.
- Built a tweet sentiment classifier with pandas text cleaning, Porter stemming and a CountVectorizer Naive Bayes model.

Peer-to-Peer Distributed Hash Table (2D#4)

Feb 2024 - Apr 2024 | Java, Bash, Azure, Wireshark

- Engineered a decentralised, fault-tolerant distributed hash table in Java and Bash with a TCP node-to-node protocol.
- Validated protocol reliability and performance on an Azure VM using Wireshark against the designed specification.

Restaurant Management System | Team Leader

Jan 2024 - Apr 2024 | Agile, Kanban, Visual Paradigm

- Led a team of 6 through Agile delivery of a restaurant management system, running the Kanban board and milestones.
- Produced the full system and workflow design in Visual Paradigm, covering use cases, class structure and process flows.

EDUCATION

King's College London | MSc Artificial Intelligence

London, UK | Oct 2025 - Sep 2026 | Dissertation submitted, awaiting final result

Modules: Artificial Intelligence Planning; Agents and Multi-Agent Systems; Agent Reasoning and Decision Making; Data Mining; Pattern Recognition; Neural Networks and Deep Learning; Computer Vision; Machine Learning.

City, University of London | BSc Computer Science, First Class Honours

London, UK | Sep 2022 - Jul 2025

- Outstanding BSc Project Prize, 2024-2025, Department of Computer Science, School of Science and Technology.
- Outstanding Stage 1 Student Award, 2022-2023, Department of Computer Science, School of Science and Technology.

Modules: Data Structures and Algorithms; Computer Networks; Cloud Technology; Introduction to Artificial Intelligence; Programming and Mathematics for Artificial Intelligence; Computer Vision; Individual Project.

ADDITIONAL INFORMATION

Languages: Kurdish (native), Persian (fluent), Turkish (fluent), English (fluent)

Certifications: Google IT Support (Coursera), Google Cloud, Flutter (Zero to Mastery), Python 3 (Codecademy), Advanced C# with WinForms and WPF