Ahmed Mahfouz

Utrecht , NL • Portfolio website • GitHub • contact@ahmedmahfouz.me • (+31) 6 3562 7674

Education

Eszterházy Károly Catholic University Eger, Hungary

Computer Science BSc. **First in class**. GPA : 4.28 September 2020 - June 2023

Relevant Coursework: Application of Neural Networks and machine learning, Data analysis, Discrete Math, Linear Algebra.

Awarded Stipendium Hungaricum full scholarship.

Thesis: AI beats a game using N.E.A.T algorithm.

Delta University for Science and Technology Damietta, Egypt

Telecommunication Engineering BSc. **Fifth in class**. GPA : 3.678 September 2018 - May 2020

Relevant Coursework: Engineering Math 4

Technical Skills

Languages & Paradigms: Python (2.x & 3.x), Java, JavaScript/TypeScript, C++, C#, PHP, SQL, HTML5 & CSS3 (SASS)

Frameworks & Libraries: React, React Native (Expo), Node.js/Express, Flask, Tornado, OpenCV, NumPy, PyTorch, TensorFlow, WhisperX, vLLM, YOLO

State & Data Management: Legend-State, Redux, SQLite, Supabase (Postgres + real-time), MongoDB, MySQL, PostgreSQL, Weaviate

DevOps & Containers: Docker, OSGi & Apache Karaf, Gradle, Nexus, CI/CD (GitHub Actions/Jenkins), Expo OTA updates

Observability & Monitoring: OpenTelemetry, Prometheus, Jaeger, Grafana, NVIDIA DCGM

Tools & Productivity: Git, Jira, Postman, Cypress, VS Code, IntelliJ, Eclipse, Jupyter, LaTeX, Markdown

Language: Arabic (fluent), English (C1), Dutch (A2), Hungarian (B1)

Professional Experience

Careweb Utrecht, Netherlands

Software Developer Nov 2024 - present

- Built and deployed new RESTful microservices in Java/Spring—as OSGi bundles on Apache Karaf—to support complex person-to-person and person-to-organization relationships, boosting data retrieval performance by 30%.
- Migrated a monolithic JSP UI into a React-based front-end, **reducing page load times by 40%**.
- Designed dynamic JSP custom tags and React components for localized NL/EN interfaces.
- Orchestrated multi-artifact release pipelines with Karaf bundle deployments, streamlining CI/CD across 15+ OSGi modules.
- Developed an AI-driven transcription and summarization pipeline using Whisper, FastAPI and vLLM
- Cut **summarization latency to under 10 seconds** and reducing transcription time for a 3 minutes audio by 35%.
- Integrated NVIDIA DCGM for GPU-accelerated **processing of up to 30 000-token prompts**, leveraging dynamic batch sizing, multi-threading and DCGM-driven monitoring.
- Engineered **multilanguage (NL-EN)** context-aware chunking and semantic embedding workflows—splitting transcripts, generating embeddings via sentence-transformers and indexing in Weaviate—to power fast, accurate question-answering.
- Built an end-to-end observability stack with OpenTelemetry, Prometheus, Jaeger and Grafana: defined custom gauges and histograms for job counts, durations and GPU utilization, and created dashboards.

VOIAR (MWLC) Utrecht, Netherlands

AI Developer project link March 2023 – Nov 2024

- Developed socket communication protocols (ipc/tcp) between the robots with an **average speed of 0.0043secs**.
- Eliminated redundant training assets from a 4-million-record MySQL database by 75% from each training session.
- Enhanced mobile controller responsiveness across 4 screen sizes, improving accessibility.
- **Achieved 98% accuracy in inference** confidence overview implementation, enhancing decision-making accuracy.
- Cut inference time from 980ms to 80ms by leveraging GPU, **boosting performance by 11.25x**.
- Maintained a MongoDB-based log system for services tracking, enabling efficient **troubleshooting for 98% of issues**.
- Improved obstacle detection input by 3%, enhancing navigation precision.
- Leveraged advanced technologies for robot fleet management, achieving **command responsiveness of 0.0035 secs.**
- Ensured 99% deployment risk reduction by conducting **six CI/CD tests**, validating system across platforms.
- **Contributed to the official YOLO repository** by publishing a TensorRT-accelerated Docker image tailored for NVIDIA Jetson edge devices, reducing average inference time by over **43%**, demonstrating a **72.3% FPS improvement** and nearly **2× images per minute**, enhancing real-time robotic vision.

IoT Research Institute Eger, Hungary

Frontend Intern April – October 2022

- Wrote 10 automated tests for the front-end using Cypress E2E.
- Reported **more than 20 problems** with suggested fixes throughout production lifecycle.
- **Cut page load speed by 5%** with a responsive web UI, improving engagement.
- Follow daily stand-up meeting in scrum and agile methodology.
- **Boosted website performance by 15%** through optimizations like code minification.
- Contributed to the enhancement of interactive JavaScript games for kids on the e-learning platform.
- Created prototypes for new features, enhancing UX/UI design.

Projects

Listimate – Smart grocery list Website Google play App store

- A cross-platform grocery-list app with offline-first SQLite caching and real-time sync via Supabase, driving 2K+ installs and 20K+ weekly social-media impressions.
- Implemented a de-duplication algorithm and drag-and-drop reordering to keep lists tidy and boost daily active usage by ~25%.
- Built secure, shareable one-time links (2 hr expiry) for multi-user collaboration without exposing private data.

Storify – AI story from photo Google play

- Turn any picture you have to a story using latest LLaVA 13B parameter model, built on top of Facebook's LLaMA2.
- Cross platform code that **works on both Android and IOS**.
- Got more than **2K generated story**.

AI beats a game using N.E.A.T algorithm project link

- Surpassed human high scores by training a N.E.A.T neural network, achieving 4x speed for **30+ minutes** of gameplay.
- **Enhanced training efficiency by 40%** with a Pythagoras theorem-based pixel-perfect collision system.
- Boosted AI vision capabilities with an **advanced 50+** line vision algorithm, improving navigation and detection.
- Implemented particle system.

Letter combination in analog phone **project link**

- Utilized DFS algorithm, **reducing runtime by 40%,** demonstrating efficiency improvement.
- Developed a Java messaging app, translating numeric to dial-pad text.
- Implemented regex-based validation, filtering invalid entries, enhancing app performance and reliability.
- **Executed 16-unit tests**, achieving **>95% code coverage**, validating across special characters and null values.

MNIST **project link**

- Implemented a handwriting image recognition algorithm with an **efficiency of 95%.**
- Worked on dealing with training data in a MySQL database.

Cyclone Game **project link**

- An Arduino game inspired by the old arcade “Cyclone” game
- Got featured on instructables website for **best circuits project**.
- Made a tutorial and got more than **1.5+ views** with people sending their creation of it.