

# Ayman Mahfuz

512-705-8897 | [ayman.afeef@gmail.com](mailto:ayman.afeef@gmail.com) | [linkedin.com/in/aymanmahfuz](https://www.linkedin.com/in/aymanmahfuz) | [aymanmahfuzportfolio.com](https://aymanmahfuzportfolio.com)

## EDUCATION

---

### The University of Texas at Austin

Austin, TX

*Bachelor of Science in Computer Science, Concentration in AI and ML*

Aug 2023 - May 2027

- **Courses:** Data Structures, Algorithms, Operating Systems, HPC, Linear Algebra, Calculus, Probability

## EXPERIENCE

---

### Platform Validation Intern (Incoming)

May 2025 - Aug 2025

*Arm*

- Selected for an AI-focused platform validation internship, targeting optimization of ARM architectures for LLM workloads.
- Internship will involve hardware-software co-design, benchmarking AI pipelines, and systems-level validation.

### Research Assistant

Jan 2025 – Present

*The University of Texas at Austin - Artificial Intelligence Lab, Texas Robotics*

- Designing multiagent RL policies from the ground up to train NAO robot soccer team, progressing from walking and dribbling to defending and attacking; blending curriculum learning, distillation, and hierarchical models across 5M+ simulation episodes for seamless real-world play.
- Accelerated a 400K+ line C++ robotics stack, boosting state-space efficiency 5x and cutting latency to sub-50ms; tuned simulation dynamics to mirror reality, enabling rapid policy design and deployment through sim-to-real pipeline.

### Data Engineer Intern

Jan 2025 – Present

*The Sunwater Institute*

- Developing production LLM system targeting daily analysis of 1000+ legislative documents, with RAG-enhanced generation achieving 85%+ human-rated accuracy on policy insights.

### Software Engineer, Research Assistant

Aug 2023 – May 2025

*The University of Texas at Austin - Center for Media Engagement*

- Built scalable ETL pipelines for 150M+ news articles, leveraging APIs, sitemaps, Pandas, and BigQuery; delivered real-time monitoring dashboards with Python and SQL.
- Fine-tuned BERT models for clickbait detection, story classification, entity recognition, and sentiment analysis, achieving up to 99% accuracy.
- Independently designed and deployed a full-stack research platform (React, Flask, Firebase) with 3 interactive games, MTurk integration, and 99.99% uptime for 1,000+ participants.

### ML Engineer, Research Assistant

Aug 2023 – Jan 2025

*The University of Texas at Austin - Oden Institute for Computational Engineering and Sciences*

- Developed a high-throughput 3D pancreas MRI segmentation pipeline using TACC's H100 GPUs, Apptainer containers, and SLURM scheduling; improved Dice score by 12%, matching state-of-the-art results.
- Benchmarked CNNs, transformers, and hybrids (e.g., PanSegNet) on 1000+ MRI scans, resolving GPU memory, I/O, and mixed precision issues to enable scalable training.

## PROJECTS

---

### Inkwell: Youtube for Books | *Django, React, PostgreSQL, Docker*

- Founded and built "Inkwell," a scalable React/Django platform for book-sharing, featuring 50+ REST endpoints, JWT auth, intelligent search, and AWS-integrated deployment.

### Leetcode Matchmaker | *Flask, React, K-Neighbors*

- Developed a web application that compares 2000+ LeetCode problems, finding and displaying similar problems using cosine-similarity on problem vectors, leveraging machine learning techniques and built with React and Flask.

## TECHNICAL SKILLS

---

**Languages:** Python, JavaScript, Java, C, C++, SQL, HTML/CSS

**Frameworks/Libraries:** React.js, Flask, Django, Pandas, PyTorch, TensorFlow

**Systems/Tools:** Google Cloud Platform (GCP), Git, AWS, HPC, CUDA