Wael Fahmy

Trustworthy Machine Learning | Uncertainty and Distribution Shift

Ontario, Canada · wfahmy@lakeheadu.ca · wael.anwar.fahmy@gmail.com

ORCID: 0009-0006-4749-7140 · Google Scholar

linkedin.com/in/wael-fahmy · github.com/wael-fahmy · waelfahmy.com

EDUCATION

MSc in Computer Science, Lakehead University, Thunder Bay, ON Sep 2024 to Sep 2025

Course-based program. GPA 4.0/4.0, **94%** overall average; A+ in every graded course.

BSc in Electrical Engineering, Cairo University, Egypt Sep 2015 to May 2020

Computer and Communication major. GPA 3.32/4.0 (WES evaluated). Graduation project grade A.

RESEARCH INTERESTS

My research focuses on the reliability of machine learning under uncertainty and distribution shift, using conformal prediction, calibration, causal inference and rigorous evaluation, with applications in robot learning, computational biology and health.

PREPRINTS

- **W. Fahmy**, E. Krikun, A. Alkhateeb. *ProtCOVER: Family-Aware Conformal Prediction for Variant Effect Predictors on Unseen Protein Families.* bioRxiv, 2026, doi:10.64898/2026.11.13.766406. Submitted to RECOMB 2027. Code: github.com/wael-fahmy/protcover.
- **W. Fahmy***, E. Krikun*, A. Alkhateeb. *Design and outcome-model estimability in doubly robust survival analysis: a breast cancer reanalysis.* medRxiv, 2026, doi:10.64898/2026.09.29.26364358. (*equal contribution)
- E. Krikun, **W. Fahmy**, A. Alkhateeb. *Evaluating the Identifiability of Temporal Delay and Response Attenuation in Endometrial Decidualization.* bioRxiv, 2026, doi:10.64898/2026.11.03.764027. Submitted to RECOMB 2027.

RESEARCH EXPERIENCE

Robot Trust Research Program | Lakehead University | **Nov 2026 to Present**

Lead researcher with Prof. Abedalrhman Alkhateeb

Designed four ongoing studies, built shared simulation and evaluation pipelines, and froze each decision rule (timestamped, SHA-256 hashed) before results were read. Results are preliminary.

- **Calibrated world models:** the base NVIDIA Cosmos3-Edge model showed no cup tipping in 224 clips. After LoRA adaptation on 2,500 simulated clips, tip rates had a mean absolute error of 0.035 against simulator rates, including held-out pushes. A per-clip reward degraded calibration; a group-level filter preserved it for one round, not two.
- **Physics-consistent augmentation:** in MuJoCo grasping with diffusion policies, glass breakage rose from 2 to 99 of 120 episodes per condition as the share of appearance-only augmentation rose from 0% to 75%. At a 50% share, physics-consistent augmentation cut breakage by 56 percentage points; a graded follow-up suggests policies copy demonstrated grips.
- **Generated 3D environments:** against laser scans, NVIDIA Lyra 2.0 geometry was reliable only within about 1 m of observed regions. Developing conformal uncertainty maps, with a sealed test on held-out rooms in progress.
- **Vision-language safety checks:** across four simulated hazard families, one misleading camera view cut a frontier model's stop rate from 0.96 to 0.43; carrying earlier observations forward restored it to 1.00.

RESEARCH EXPERIENCE (CONTINUED)

Machine Learning for Biology and Health | Lakehead University | **2026 to Present**
Research collaborator with Prof. Abedalrhman Alkhateeb and Elena Krikun

- **ProtCOVER:** developed family-level conformal calibration for protein variant effect predictors to quantify uncertainty on unseen protein families. Studied coverage under stated assumptions and their limitations; released the implementation, tests and reproducibility scripts.
- **Survival analysis:** developed a diagnostic for outcome-model estimability prior to doubly robust survival estimation; reanalysed TCGA-BRCA data from 1,076 patients.
- **Endometrial decidualization:** contributed analysis, writing and manuscript review to an identifiability study of temporal delay and response attenuation parameters.

MANUSCRIPTS IN PREPARATION

Robot Trust program (drafts written; not peer reviewed):

- **W. Fahmy**, A. Alkhateeb. *Calibrated World Models: Do Fine-Tuned Video World Models Show Physical Outcomes at the Right Odds?*
- **W. Fahmy**, A. Alkhateeb. *Augmentation That Lies About Physics.*
- **W. Fahmy**, A. Alkhateeb. *Which Parts Are Real? Certified Uncertainty for Generated 3D Worlds.*
- **W. Fahmy**, A. Alkhateeb. *What the First Look Knew: Vision-Language Safety Checks for Robot Arms under Blurred and Doctored Camera Views.*

Other: **W. Fahmy**. *Malware image classification with EfficientNet on MaleVis and Malimg* (99.2% accuracy; MSc project), in preparation for an IEEE journal.

RESEARCH SOFTWARE AND COMPUTING

- Run training and simulation experiments (MuJoCo, NVIDIA Cosmos) with Slurm on the Digital Research Alliance of Canada's H100 clusters.

SELECTED PROFESSIONAL EXPERIENCE

Charti Health AI Inc. | Thunder Bay, ON (Remote) | **Feb 2026 to Present**
Founding Engineer & Technical Lead (part-time since Jun 2026)
Sole engineer on a clinical AI system used in Canadian clinics: real-time speech processing, self-hosted LLM inference with vLLM, about 2,100 automated tests and distributed tracing.

Fitlog | Montreal, QC (Remote) | **Dec 2022 to Present**
Senior Full Stack Engineer, AI & Cybersecurity (since Sep 2025)
Built a Python AI gateway with model routing, validated outputs, fallbacks and 391 tests; develop health data applications and LLM features for a platform used by 27,500+ coaches and 350,000+ clients.
Other roles (2020 to 2024): software engineering in education, banking and document automation.

RESEARCH SKILLS

Methods: Conformal prediction, calibration, causal inference, survival analysis, identifiability, pre-specified evaluation

Tools: Python, PyTorch, Hugging Face, LoRA, MuJoCo, NVIDIA Cosmos, LeRobot, Slurm, Docker, Git, LaTeX

Other MSc projects: hyperspectral imaging (GAN augmentation), driver drowsiness detection, energy demand forecasting, LLM coding agent comparison

Languages: English (IELTS L 8.0, R 8.5, S 7.0, W 7.0), Arabic (native)