

Jack Warren

jackwarren430@gmail.com | 4044316150 | www.linkedin.com/in/jack-warren-89605a135 | https://github.com/jackwarren430

PROFESSIONAL SUMMARY

ML/AI Engineer and CS student at Georgia Tech (B.S. 2026, M.S. 2028) with hands-on experience fine-tuning and deploying LLMs, building ML research pipelines, and shipping production software. Raised model accuracy from 75% to 96% via QLoRA fine-tuning across an 8-billion-transaction dataset at Equifax; improved physics-informed neural network MSE by 12% at GTRI. Fluent across the ML stack from model internals and quantization to HPC job scheduling and agentic system design.

EDUCATION

Georgia Institute of Technology

B.S. Computer Science, Minor in Japanese / M.S. Computer Science (expected May 2028).

May 2026

GPA: 3.77 / 4.0.

Relevant Coursework: Deep Learning, Machine Learning, Probability & Statistics, Knowledge-Based AI, Algorithm Analysis, Computer Systems

EXPERIENCE

Data & Analytics Intern

May 2026 – August 2026

Equifax

Atlanta, GA

- Fine-tuned and evaluated Gemma 3 4B and Gemma 4 E4B via QLoRA on Vertex AI Workbench to classify ~8B financial transactions across 30 primary and ~300 total categories
- Adapted Gemma 4 E4B for legacy infrastructure by modifying model internals to support older library versions
- Tuned quantization and dtype configurations for NVIDIA T4 GPUs, raising test accuracy from 75% to 96%

Student Assistant Researcher

January 2026 – May 2026

Georgia Tech Research Institute

Atlanta, GA

- Extended a physics-informed neural network surrogate predicting intensity signal sequences from input physical parameters
- Designed and A/B-tested an intermediary value loss module anchoring the physics loss to mid-network intermediate predictions
- Implemented push-forward training, feeding model predictions back as inputs during training, improving test MSE by over 12%
- Built a YOLO-based object-detection pipeline for moving vehicles using GPS coordinates to assign global position labels

Team Lead, Automated Algorithm Design Researcher

August 2024 – December 2025

Georgia Tech VIP

Atlanta, GA

- Applied Multi-Objective Genetic Programming with LLMs to evolve a CNN image classifier; benchmarked single-LLM and multi-LLM island configurations to evaluate population diversity tradeoffs
- Designed a multi-LLM island architecture supporting multi-population evolution, migration, and checkpointing
- Built a reproducible HPC pipeline on PACE-ICE using SLURM sbatch, H100/A100 allocation, and shared model storage

Software Developer Intern

May 2024 – August 2024

MessageGears

Atlanta, GA

- Standardized Angular/JSP actions UX with confirmation modals and shipped multiple UI fixes
- Rewrote Groovy/Spock API tests and split suites by endpoint, cutting runtime from minutes to seconds and accelerating builds

PROJECTS

Spectral Text Encoder

- Built a spectral text encoder mapping tokens to sine triples summed into a fixed-length waveform embedding
- Raised in-batch contrastive validation accuracy from 56% to 78% using frequency-separation and per-channel InfoNCE losses

Wave Scalar Field Sequence Generator

- Built a PyTorch video transformer learning 2D wave dynamics from height frames via block-causal attention and RoPE
- Reduced rollout error via a pushforward horizon schedule, velocity-supervision aux head, and residual conv refinement

Agent-Powered Creator-Game

- Built an open-ended game converting spoken ideas into explorable 3D scenes using an agentic Gemma 4 loop with custom tools
- Ran fully local inference on Apple Silicon using quantized models, FLUX for image generation, and TRELIS for 3D mesh

Self-Driving Car Algorithm

- Implemented TransFuser to fuse image and LiDAR representations via cross-attention and convolutions, trained on CARLA

SKILLS

Languages: Python, SQL, C, Java, JS/TS, Groovy

Libraries: PyTorch, Transformers, TRL, PEFT, Unsloth, DEAP/GP, Sklearn, Pandas, NumPy, OpenCV, Streamlit

Technologies: CUDA, Hugging Face, SLURM (HPC), Google Cloud Platform, Vertex AI Workbench, BigQuery, AWS S3