

Nicholas Hoffs

(949)-280-2672 | niethoffs@gmail.com | nichoffs.com | github.com/nichoffs

EDUCATION

UNIVERSITY OF VIRGINIA

August 2022 - Present

Computer Science B.S / Mathematics B.A. - 3.9 GPA

- Coursework: Data Structures & Algorithms, Computer Architecture, Math for Data Science, Software Development, Discrete Mathematics, Physics, Chemistry, Linear Algebra
- Clubs & Activities: Assistant-Coach/Player UVA Men's Club Water Polo, Gracie Jiu-Jitsu

EXPERIENCE

University of Virginia

August 2024 - Present, Charlottesville, VA

PhD Research Assistant

- Integrating multimodal knowledge graphs (MMKG) with **large language models** (LLMs) for Question Answering (QA), utilizing an innovative **graph neural network** aggregation and cross-modal attention algorithm.
- Developed and curated novel MMKG QA dataset with over 200k samples for training graph model.

Cavalier Autonomous Racing

May 2024 - Present, Charlottesville, VA

Perception Team Member

- Implemented and trained NVRadarNet model to improve opponent bounding box detection accuracy (+22% IOU) compared to previous approach.
- Optimizing latency (<20ms) and reliability across perception stack with **ROS2** in both C++ and Python.

Caju AI

May 2024 - August 2024

Staff AI Engineer

- Designed and implemented comprehensive evaluation suite for all AI pipelines, including RAG (Retrieval-Augmented Generation), sentiment classification, and entity extraction systems.

Unbox

June 2020 - February 2021

Website Developer and Designer

- Designed and developed website using HTML, CSS, and JavaScript to interface with GIS API, delivering critical public food infrastructure information to families in need; supported thousands of active users.

PROJECTS

Implemented deep learning library from scratch in Python/NumPy and pure C

- Created custom tensor automatic differentiation library from scratch in pure C and Python/NumPy.
- Executed training runs on MNIST image data, achieving **95% accuracy** for C and Python implementations.

Reproduced Mechanistic Interpretability paper in TinyGrad

- Reproduced "Progress measures for grokking via mechanistic interpretability" paper using TinyGrad, achieving perfect generalization accuracy of 100% on test set based on modular addition task.
- Confirmed reproduction using FFT analysis to verify model's learned periodic behavior.

Trained custom LLM from scratch

- Implemented and trained **125M-parameter GPT model** on TinyStories dataset using **TinyGrad** library, allowing for scalable training using a single GPU or multi-GPU through distributed data parallel (DDP).
- Generated ~coherent English stories using GPT, achieving train loss of 1.9 over 24-hour training period.

TECHNICAL SKILLS

- Machine Learning: PyTorch, TinyGrad, Pandas, LangChain, OpenAI API, FastAPI, PostgreSQL, pgVector
- DevOps: Git, GitHub, GitLab, Docker, Jira, AWS Lambda, AWS Step Functions, AWS EC2
- Programming Languages: Python, C/C++, Javascript, TypeScript, Java, MATLAB, SQL, HTML, CSS
- Robotics Software/Hardware: ROS2, relpy, relcpp, LIDAR, RADAR, Camera