Vihanga Mihiranga Malaviarachchi

+61 0438485711 | vihaaanga.mihiranga@gmail.com | github.com/VihangaFTW | linkedin.com/in/vihanga-malaviarachchi

Education

Monash University **Graduated: December 2025**
Bachelor of Computer Science — Algorithms and Data Structures *Australia*

Monash University **Expected Completion: December 2026**
Bachelor of Computer Science (Honours) — Deepfake Audio Detection *Australia*

Technical Skills

Languages: JavaScript, TypeScript, Python, Go, Rust, C, Java

Developer Tools: AWS (ECR, EC2, RDS, EKS), Docker, Git, GitHub Actions, Vite, Postman

Libraries/Frameworks: React, TanStack Start, Tailwind CSS, Gin, gRPC, LangChain, LangGraph, ChromaDb, PyTorch

Projects

ByteTok: BPE Tokenizer for NLP Research | *Source Code* **Python | Rust | NLP | PyPI**

- Engineered a byte-level BPE tokenizer with a Rust-based training algorithm that rivals HuggingFace's BPE tokenizer in performance, with **versioned** `.model/.vocab` **serialization** and **automatic tokenizer detection** on load.
- Implemented a Rayon-based **parallel processing pipeline** with GIL release, enabling a mininmum of **15 MB/s encoding** and **77M tokens/s decoding** throughput for large batches, with support for **special tokens.**
- Built production packaging and **release automation** with cross-platform wheels (Linux/macOS/Windows) and PyPI publishing via GitHub Actions.

LiminalGPT: Prose Emulator | *Source Code* **Python | Deep Learning | PyTorch**

- Built a **decoder-only transformer** to emulate surrealist style prose, utilizing *ByteTok* tokenization for nuanced text generation.
- Curated and pre-processed a custom dataset of personal creative writing excerpts to train the model on specific syntactic patterns and thematic elements.
- Implemented **pre-norm** architecture with **multi-head self-attention** and **causal masking**, achieving optimal text quality at **1,500** training iterations.

Monex: Distributed Financial Backend | *Source Code* **Go | gRPC | PostgreSQL | AWS EKS**

- Engineered a comprehensive banking backend in Go, utilizing PostgreSQL and SQLC to manage **atomic money transfers** across multi-currency accounts.
- Designed a dual-protocol interface (gRPC + REST via gRPC-Gateway) with structured **observability** for request tracking and **email verification**.
- Implemented secure **session management** and automated **CI/CD pipelines** for **containerized deployment**.

Autodiff: Neural Engine | *Source Code* **Python | Deep Learning | Autograd**

- Engineered a minimal autograd engine for scalar computations from scratch, implementing backpropagation and **automatic differentiation** via a dynamic computational graph.
- Developed foundational deep learning components including Neurons, Layers, and Multi-Layer Perceptrons (MLP) with full gradient descent support.

AI Agent Development | *Source Code* **LangGraph | LangChain | ChromaDB**

- Developed a workflow assistant using LangGraph, featuring an advanced **RAG** pipeline with **HyDE** and **Multi-Query Expansion**, orchestrated by a third-party LLM.
- Integrated ChromaDB for semantic retrieval to enhance the accuracy and context-awareness of generative responses.

EduTaskSync | *Source Code* **React | TypeScript | Tailwind CSS | TanStack Query**

- Engineered a collaborative task management platform to centralize fragmented student workflows, replacing high overhead enterprise tools like Jira with a semester focused React Kanban system.
- Implemented real-time **state synchronization** and **optimistic updates** using TanStack Query to ensure seamless coordination across distributed Monash student teams.