

```
~/lee $ whoami --high-signal
```

Lee Penkman

Machine learning engineer, VFX tools builder, and product founder working across **generative AI, graphics, infrastructure**, and **developer-facing systems**. I like shipping useful software: products that people can try, repos that other builders can fork, and pipelines that survive real production load.

leepenkman@gmail.com

lee.app.nz

github.com/lee101

AI / ML

VFX / Graphics

Systems / Product

CURRENT VECTOR

Building an ecosystem of applied AI products: model gateways, local agents, AI media tools, accelerated inference, semantic memory, speech control, and specialized product experiments.

PRODUCT SURFACE

More product weight, less chronology. These are the projects I would lead with for AI, infra, graphics, and applied research roles.

AI CODING PLATFORM **Codex Infinity**



Autonomous coding-agent platform with cloud fleet orchestration, GitHub integration, streamed sessions, CI repair loops, and multi-model support.

codex-infinity.com

MODEL GATEWAY **OpenPaths**



OpenAI-compatible gateway across chat, image, video, music, speech, transcription, and embeddings, with provider routing and reasoning controls.

openpaths.io [repo](#)

ACCELERATION **CuteDSL**



Custom CUDA/Triton kernels, fused model pipelines, Chronos-2 acceleration, and an API surface for fast AI inference and training.

cutedsl.app.nz cutedsl.cc [repo](#)

AI MEDIA/COMMUNITY **Netwrck**



Social AI community with chatbots, stories, generated art, voices, and agent-style experiences.

netwrck.com [stable-diffusion-server](https://stable-diffusion-server.com)

UNIFIED AI API **Text Generator**



Text, speech, vision, code generation, prompt tooling, OCR/caption enrichment, and self-hostable model-serving infrastructure.

text-generator.io [repo](#)

SPEECH CONTROL **DictatorFlow**



Desktop and browser dictation: hotkey recording, local GPU Parakeet ONNX transcription, cloud fallback, command mode, and transcript history.

dictatorflow.com [repo](#)

3D TOOLS **SimplexGen**



3D generation and editing app with a Go server, Vite frontend, local asset indexing, mesh tools, generation queues, billing, and R2 storage.

simplexgen.com

GAME **Ethereal Storm**



Go/Ebiten twin-stick bullet hell with co-op, procedural terrain, wave survival, bosses, spell combinations, progression, and dynamic audio.

etherealstorm.app.nz

HIRING / MARKETPLACE

Hires.nz



Recruitment and product experiments in the same workspace as SimplexGen, with deployed web assets and Go-backed services.

hires.nz

VOICE/PRODUCT COMMS

Ring.nz



Modern web communication product with auth, checkout, chat/call UI, Go backend checks, and frontend verification workflows.

ring.nz [repo](#)

AI MEMORY

GoBrain

GB

Hybrid semantic and keyword memory store for agents, with 150K+/sec CPU embeddings, MCP tools, knowledge graph links, and CLI workflows.

gobrain.app.nz [repo](#)

FORECASTING

Stock Prediction

RL

Auto-research and RL trading experiments over Alpaca/Binance, selector evaluation, Chronos-style forecasting, and realistic portfolio probes.

[repo](#)

AI DATA ANALYST

Helix

HX

AI data scientist product for interpreting datasets, generating insights, and automating reporting with model-backed analysis workflows.

helix.app.nz

EXPERIMENT OPTIMIZATION

OptimizeFlow



Optimization tooling with neural bandits, tracking, metrics, backend SDKs, personalization, and experiment-flow service code.

optimizeflow.com [brand](#) / [local codebase](#)

OPEN SOURCE

Repos that show the implementation style: GPU serving, API compatibility, local-first tools, Go services, Python ML systems, and frontends.

Stable Diffusion Server

Production-ready image generation server with Flux/SDXL, ControlNet, inpainting, R2/GCS storage, FastAPI, and Gradio tooling.

[netwrck/stable-diffusion-server](#)

CuteDSL

AI kernel and model acceleration work spanning Triton, CUDA, torch.compile, and equivalence-focused benchmarking.

[lee101/cutedsl](#)

Text Generator

Self-hostable API-compatible text, speech, vision, OCR, prompt, and model-routing stack.

[lee101/text-generator.io](#)

GoBrain

Agent memory with Postgres, semantic search, BM25, RRF fusion, MCP server, REST API, and CLI.

[lee101/gobrain](#)

Ring.nz

Web communications app with Go backend, Vite frontend, deployment guardrails, and visual verification.

[lee101/ringnz](#)

Stock Prediction

RL and forecasting research code for market prediction, simulations, selectors, backtests, and agentic training runs.

[lee101/stock-prediction](#)

EXPERIENCE

Condensed chronology from the old CV, tuned to keep attention on technical responsibility and product impact.

2025 - Present

Founder / Builder - Applied AI Product Ecosystem

Independent product and open-source work across AI gateways, coding agents, speech tooling, model acceleration, media generation, 3D tools, games, and forecasting.

- Shipping productized AI infrastructure rather than demos: billing, auth, API docs, queues, monitoring, provider routing, and deployment workflows.
- Mix of Go, Python, CUDA/Triton, TypeScript/React, FastAPI, Postgres, R2/GCS, and local GPU pipelines.

Aug 2023 - Aug 2025

Senior Machine Learning Engineer - Microsoft Clipchamp

Worked on video creation/editing ML features in collaboration with research teams, bringing AI features into a production creative product.

- AI video creation, AI video editing, background removal, AI VFX, NLP, PyTorch, Hugging Face, and ffmpeg-heavy workflows.

Aug 2021 - Jul 2022

Founding Machine Learning Engineer - Compose.ai YC W21

First ML engineer for an AI writing assistant, improving product quality and the model/data processes behind it.

- Autocomplete, rewriting, spell checking, dataset curation, bias reduction, chatbot/assistant work, training pipelines, and LLM inference acceleration.

May 2019 - Aug 2021

Image Processing / Machine Learning Engineer - Canva

Delivered image processing and ML capabilities for photo editing, after earlier security engineering work.

- Background removal, object detection, training data pipelines, GLSL/OpenGL, PyTorch, Terraform, AWS, Docker, phishing detection, incident response, and CI/CD security.

Sep 2015 - May 2018

Motion R&D Technical Director - Weta Digital

Supported animation, motion editing, crowd simulation, mocap, and virtual artists with production tools and reliability engineering.

- Scheduling, statistics, ML/big-data analytics, data warehousing, artist workflow tools, image processing, camera warping, publishing/versioning, and distributed 3D data pipelines.

2010 - 2015

Earlier Engineering - Fairfax, .nz Registry, Infoxchange, 99designs, Kiwiplan, University of Auckland

Tech lead, developer, intern, and teaching assistant roles across high-traffic media, DNS registry analytics, government service maps, ecommerce, BI, and CS teaching.

- AWS, Django, Rails, Python, PHP, Go, JavaScript, Hadoop/Hive/HBase, Scala/Spark, Postgres, Docker, D3, ETL, and test automation.

ACADEMICS

Collapsed by default: the headline is that the academic record backs up the systems, algorithms, graphics, and teaching thread.

BSc Computer Science - University of Auckland

- Completed Computer Science in two and a half years.
- A/A+ dominant record across programming, algorithms, distributed systems, AI, graphics, operating systems, databases, and computational science.
- Teaching Assistant for Java, Python, and practical computing; genuinely enjoys teaching people how computers work.
- Top in class for COMPSCI 369 Computational Science; J.C. Butcher Award in Theoretical Computer Science; Kiwiplan Scholarship for excellence in CS/Software Engineering.
- First place in CS101 game programming; first place in CS320 fastest C quickselect; 6th in the 2012 New Zealand Programming Competition tertiary open.

► **Expand selected academic results**

SIGNALS

Extra credibility markers that are useful in interviews and portfolio conversations.

Patent

AI ensembles for image background removal.

patents.google.com

Technical Range

- PyTorch, Hugging Face, CUDA, Triton, ONNX Runtime
- Go, Python, TypeScript, React, FastAPI, Django
- Postgres, R2/GCS, AWS/GCP, Docker, Terraform

Creative Range

- Image/video generation and editing
- VFX, motion, mocap, crowd simulation tools
- Games, 3D assets, graphics pipelines, procedural systems

References available on request. Contact: leepenkman@gmail.com. Original long-form CV source: lee.app.nz.

Print layout is intentionally lighter: use browser print to save a clean PDF.