

DIEGO MAGANA

Los Angeles, CA • diegomagana@gmail.com • (480) 544-2890 • github.com/diego-magana • diego-magana.github.io

Principal engineer with ~10 years across data science, ML, and engineering leadership, now pursuing an MSCS specializing in machine learning. Owns enterprise cloud, data, and ML infrastructure end-to-end, and has built modern deep-learning systems from first principles — autograd, language models, and GPT-2 pretraining — including an original interpretability analysis of learned representations. Focused on mechanistic interpretability and alignment.

EDUCATION

Johns Hopkins University

Baltimore, MD

Master of Science, Computer Science (AI) — in progress · GPA 3.93

Amherst College

Amherst, MA

Bachelor of Arts, Environmental Studies (2017)

- NCAA Division III Men's Basketball · NESCAC Champions (2014) · NCAA Semifinalist (2014)

ML PROJECTS

micrograd — Autograd Engine from First Principles · github.com/diego-magana/micrograd

- Built a scalar-valued reverse-mode autodiff engine from scratch — dynamic graph via operator overloading — with gradients verified against PyTorch to 1e-5 across all operators; trained an MLP (337 params) to 100% accuracy on `make_moons`.

makemore — Character-Level Language Model Series · github.com/diego-magana/makemore

- Implemented the full character-level progression — bigram → MLP (Bengio 2003) → WaveNet (77K params, 2.01 validation NLL) — including a from-scratch backprop variant matching PyTorch autograd to within ~0.001 nats.
- Authored an original probing analysis of the trained model's learned representations — embedding geometry, gradient-based context sensitivity, and single-unit ablation — with explicit reporting of correlational vs. causal limits.

nanoGPT — GPT-2 (124M) Pretraining from Scratch

- Implemented the full GPT-2 pretraining pipeline from scratch in PyTorch on a single RTX 5080 — tokenization, transformer architecture, training loop, and evaluation.

Production ML Infrastructure · Docker · GitHub Actions · Modal · W&B

- Established a reproducible ML engineering stack across the portfolio: containerized environments, CI/CD (automated test and lint gates), a Modal inference endpoint, and W&B experiment tracking on every run.

EXPERIENCE

Raxwell Inc., Irvine, CA

Jan '23 – Jun '26

Principal Engineer (Contract)

Oct '25 – Jun '26

- Sole technical owner of the company's full Microsoft stack — Azure, M365, and Business Central — responsible for architecture, integration, and production operations.
- Deployed production LLM capabilities on Azure AI Foundry — fine-tuned models and agent workflows with automated quality benchmarks — and built the data integrations and automation across the stack.

Chief Information Officer

Mar '24 – Oct '25

- Owned end-to-end technology strategy and architecture through a growth phase: Azure cloud, enterprise data infrastructure, and the M365 ecosystem.

Principal Engineer (Contract)

Jan '23 – Mar '24

- Architected and led the company's migration to Azure; built the initial data platform and ML/analytics infrastructure still in production today.
- Developed ETL pipelines (Azure Data Factory), REST API integrations, and Power BI forecasting dashboards adopted as the primary operational decision tooling.

Mavco LLC, Scottsdale, AZ

Aug '17 – Jan '23

Senior Data Scientist

Aug '20 – Jan '23

- Established the company's first analytics and ML function, scaling the modeled product catalog from several hundred SKUs to tens of thousands through automated ML pipelines.
- Owned the full ML lifecycle — data engineering, feature development, model training, and production deployment of forecasting, pricing, and classification systems — contributing to a ~10-point lift in adjusted gross margin and informing seven-figure capital allocation.

Data Scientist

Aug '17 – Aug '20

- Built ML pipelines for sales and pricing prediction and developed financial models informing executive strategy.

TECHNICAL SKILLS

- **Deep Learning:** PyTorch, transformer architectures, LLM pretraining & fine-tuning, language modeling
- **Machine Learning:** scikit-learn, time-series forecasting, classification, statistical modeling
- **Infrastructure & MLOps:** Docker, GitHub Actions (CI/CD), Modal, W&B, Azure AI Foundry, Azure Data Factory
- **Languages:** Python, Java, SQL, C++
- **Research Interests:** mechanistic interpretability, AI alignment & safety, LLM evaluation