

Johnny Dunn

| github.com/jddunn | johnny.mov | linkedin.com/in/jdfive | johnnyddunn@gmail.com

Summary

Full-stack engineer in end-to-end applied AI/ML (natural language processing / computer vision): RAG pipelines, LLM fine-tuning, MCPs, deep learning, GPU-accelerated training, quantization, embedded development / edge inference, evaluation harnesses, IVRs.

Professional Experience

Machine Learning Engineer

Independent Consulting - Contract | Remote | May 2019 – March 2026

- Built edge AI surveillance on embedded SBC (3 NPUs) in Go for Aye Eye Global (3+ months). Multi-cam RTSP / WebRTC, websockets, SSE, YOLOv8n, face recognition, quantized LLMs for low-latency voice interface, shipping 5k+ factory units.
- Built unified ML pipeline in Python for Linkby (6-month project) automating feature engineering and model training from SQL for advertising data, enabling parallel training of RF, SVM, Naive Bayes, and LR models with A/B testing framework.
- Improved AI text detection accuracy from 55% to 71% for Grapple Research (5-month project), finetuned GPT-3.5, deployed with Docker, Supabase on AWS. Built scrapers augmenting training, reduced costs by 64% with document summarization.
- Developed semantic similarity matching tool for ImmenseX (stealth SaaS, 4-month work), on Weaviate / Neo4j RAG (hybrid vector / graph search), on GCP/ Linux, with vLLM serving open weights models. Reduced prompt writing time by 60%.
- Achieved 77% sentiment analysis accuracy (random forest) trained on 10K texts for HereAfterLegacy.ai; built MVP with scrapers for TV scripts, recreated digital persona of characters, deployed on Google Cloud Platform (14-month contract).
- Trained neural style transfer models with REST API exhibit applying film-inspired filters to selfies for Edelman (3-month work).
- Created user research, animated wireframes in Figma, and functional demo for MVP at Notebrush (15-month engagement).

Senior Blockchain Engineer

Tilting Point | New York, NY | Nov 2021 – Feb 2023

- Created ERC-20/721 upgradable / proxy smart contracts, built REST and gRPC APIs for NFT marketplace with SSO and containerization, and Merkle proof airdropping tool. Rescued \$1M+ in MATIC by reverse engineering Gnosis safe methods.
- Architected Go microservices off and on-chain events supporting Unity SDKs in 4+ games with user analytics at 1000+ TPS.

Software Engineer 2

eBay Inc. | New York, NY | Jun 2020 – Apr 2021

- Built REST APIs in Scala / Spring-Boot showing new ads in real-time from clicking selected item options (in under 1s latency).
- Handled E2E ML Ops infrastructure (CI/CD with K8s, Jenkins); enabled teams to deploy models to sandboxes exposing data science driven statistics. Accelerated deep learning model iterations by 50%; built scalable scrapers verifying data integrity.

Cognitive Implementation Engineer

Amelia (Formerly IPsoft Inc.) | New York, NY | Jun 2017 – Dec 2017

- Built AI assistant enterprise modules (PII redaction for HIPAA, SOC 2 compliance, Agile practices) for F500 medical suppliers.

Personal Projects

AgentOS – Agent [orchestration](#). Streaming [guardrails \(PII redactor\)](#), [graph multimodal RAG](#), [IVR](#), [HITL](#). github.com/framersai/agentos

Quarry – Self-organizing notes & QA via BERT semantic vector search and offline Ollama. github.com/framersai/quarry

tenets – MCP server for coding agents to aggregate context for prompts with neural ranking, fully offline. github.com/jddunn/tenets

Restless - Hierarchical attention network model for anti-malware; 93% accuracy, [ClAMP](#) dataset. github.com/jddunn/restless

Education

Bachelor of Fine Arts in Game Design (May 2017) – Parsons School of Design

Awards

5th Place, AT&T National VR and AR Challenge, with THiNKIMMERSIVE team. (<http://npe100.com/> / mobile app), 2017.

Skills

Programming Languages: Python, TypeScript, JavaScript, Go, Scala, Java, C#, SQL, Solidity, Bash

Frontend: React, Next.js, Vue, Nuxt, Redux, Tailwind, React Native, D3.js, three.js

Backend: FastAPI, Django, Node.js, Express, NestJS, gRPC, GraphQL, RabbitMQ, Apache Kafka, Stripe

AI / ML: PyTorch, TensorFlow, Keras, CUDA, Hugging Face, Transformers (BERT, Llama, Mistral, Whisper), PEFT, LoRA, OpenAI, Anthropic, LangGraph, LlamaIndex, DSPy, Mastra, RAGAS, DeepEval, scikit-learn, XGBoost, spaCy, MLflow, Stable Diffusion

Databases / DevOps: PostgreSQL, Supabase, DynamoDB, MongoDB, Redis, Pinecone, Weaviate, Neo4j, Elasticsearch, SQLAlchemy, Apache Spark, Airflow, Snowflake, AWS (Lambda, EC2, EBS, Bedrock), GCP, Docker, Kubernetes, Terraform

Tools: Unity / Unity3D, Jenkins, GitHub Actions, Adobe Creative, Electron, Capacitor, Android Studio, Selenium, Puppeteer