CHRISTOPHER MEJO Ph.D.

AI Research Scientist — LLM & Reinforcement Learning Systems

work@cmejo.com • (917) 747-8865 • New York, NY • cmejo.com • github.com/cmejo • linkedin.com/in/cmejo

SUMMARY

AI research scientist and engineer with 15+ years taking machine learning systems from research prototype to production, holding a Ph.D. in Physics from Harvard (quantum information and complexity theory, completed at 21) and postdoctoral training at MIT under Scott Aaronson. Core focus: LLM systems — retrieval-augmented generation, fine-tuning, agentic reasoning — and reinforcement learning for sequential decision-making. Solo-architected AI Scholar, an open-source RAG platform indexing 15M+ research papers (Google Summer of Code Fellow, 2024 and 2025). Shipped production LLM and RL systems in finance and a federated computer-vision system in healthcare. Comfortable owning systems end-to-end in small teams.

PROFESSIONAL EXPERIENCE

Primary Employment

Senior Data Scientist — Braverock LLC *(Nov 2021–Present)*

- Architected a reinforcement-learning trading system (deep Q-networks, actor-critic methods) for multi-asset portfolio optimization, improving Sharpe ratio 23% over prior strategy on microsecond-latency data, validated across both backtested and live trading.
- Built transformer-based market-prediction models for financial time series, deployed in a pipeline processing 1B+ quote-data points/day at microsecond-scale latency.
- Built a RAG system over 10M+ financial documents (SEC filings, earnings calls, research reports) for signal extraction, using custom embeddings and a production vector database.
- Directed distributed pretraining and instruction fine-tuning pipelines for domain-specific 7B–14B parameter foundation models (DeepSpeed ZeRO-3, custom CUDA memory-mapped tokenizers), reducing per-step wall-clock latency 38%.
- Implemented dynamic Heterogeneous Graph Attention Networks (HAN) modeling multi-asset cross-correlations across tick, intraday, and multi-week horizons, improving regime-shift detection accuracy 19%.
- Designed and deployed an enterprise crypto basis-yield strategy delivering an out-of-sample, beta-neutral Sharpe ratio of 2.63 on $60M+ in institutional AUM.
- Developed proprietary quantitative strategies achieving a 6.2 Sharpe ratio (Sortino 8.4, max drawdown <1.9%) over a continuous 24-month live deployment window, executing billions of dollars in trading volume across holding periods from sub-second to multi-month.
- Strategy scope spanned quantitative multi-exchange, multi-asset market making; treasury management; regime-aware risk-managed momentum; and multi- and single-asset statistical arbitrage / structural mean-reversion market-neutral strategies.

Senior Data Scientist — Cushion AI *(Apr 2021–Oct 2021)*

- Built CushionNLP, a proprietary BERT-based transformer system for bank-fee negotiation, driving a 30% increase in successful negotiations and $2M+ in additional annual customer savings.
- Architected a transaction-categorization system (ensemble methods, graph neural networks) achieving 95% accuracy across 80M+ transactions from ~300 financial institutions, cutting processing time 50% versus the prior rule-based (non-ML) system.

Research Data Scientist — RapidRads AI *(2019–Feb 2021)*

- Built a computer-vision system (EfficientNet-B7 backbone) for COVID-19 detection from chest X-rays, achieving 0.972 ROC-AUC with 96.5% sensitivity and 98.2% specificity (1.8% false-positive rate, 3.5% false-negative rate) under differential privacy ($\varepsilon = 1.2$, $\delta = 10^{-5}$) across a federated-learning deployment spanning 4 medical institutions.
- Combined multiple CNN architectures via ensemble methods with Grad-CAM visualization and uncertainty quantification for clinical validation.

Senior Data Scientist & Cybersecurity Researcher — Confiant *(2017–2019)*

- Built threat-detection models (Isolation Forest, One-Class SVM, LSTM autoencoders) and GraphSAGE-based botnet-detection models identifying coordinated attacks across 1B+ daily ad requests.

CEO & Lead Architect — Cronode *(2011–2017)*

- Built a recommendation engine (TF-IDF, cosine similarity, collaborative filtering) and custom Solr ranking configuration.

Concurrent Research & Fellowships

Visiting Research Scientist (part-time) — Sandia National Laboratories — Quantum Computing & High Energy Physics, Joint Google Quantum AI Initiative *(2020–2024)*

- Senior researcher in a joint research initiative with Google Quantum AI developing variational quantum eigensolvers for lattice gauge theory calculations.
- Designed quantum error correction protocols (surface codes) and resource-estimation frameworks for near-term quantum devices (Qiskit, Cirq).
- Ran distributed quantum many-body simulations (50+ qubits) on SLURM-managed clusters; contributed to DARPA and IARPA research proposals.
- Active Top Secret/Sensitive Compartmented Information (TS/SCI) security clearance.

Google Summer of Code Fellow, 2024 & 2025 — AI Scholar Project — R Foundation for Statistical Computing *(2024–2025)*

- Architected AI Scholar (github.com/cmejo/AI_Scholar), combining dense (FAISS), sparse (BM25), and graph-based (Neo4j) retrieval for scientific literature search across 15M+ research papers.
- Built an adaptive retrieval system using meta-learning to personalize search strategy, improving relevance 40% (NDCG@10) over baseline on a 200-query held-out set with citation-derived relevance judgments.
- Implemented chain-of-thought reasoning over the RAG pipeline with explicit source attribution and calibrated uncertainty (Monte Carlo dropout, evidential deep learning).
- Built a dynamic knowledge-graph construction pipeline using named entity recognition (spaCy, BERT-NER) and relation extraction to connect concepts, authors, and methodologies.

EDUCATION

Postdoctoral Fellowship, Theoretical Computer Science — Massachusetts Institute of Technology (2010–2012). Faculty Advisor: Dr. Scott Aaronson.

Ph.D. in Physics — Harvard University (2007–2010). Quantum information and complexity theory; completed at 21.

B.S. in Physics and Computer Science — Harvard University (2004–2007).

TECHNICAL SKILLS

Languages: Python, C, C++, R, Assembly, Java, JavaScript

ML / DL: PyTorch, TensorFlow, JAX, Transformers, Reinforcement Learning (DQN, actor-critic, MAML)

LLM / RAG: Retrieval-augmented generation (dense, sparse, and graph-based retrieval), FAISS, BM25, Neo4j knowledge graphs, custom embedding models, vector databases, LangChain, LangGraph, vLLM, named entity recognition (spaCy, BERT-NER), chain-of-thought and agentic reasoning, uncertainty quantification (Monte Carlo dropout, evidential deep learning), retrieval evaluation (NDCG, MRR)

MLOps: MLflow, Kubeflow, Apache Airflow, Docker, Kubernetes, Ray, Weights & Biases

Quantum Computing: Qiskit, Cirq

Distributed / HPC: SLURM, multi-GPU distributed training