Jack Lutz

(516) 509-2108 | jwlutz65@gmail.com | linkedin.com/in/jwlutz | github.com/jwlutz | jwlutz.com

Education

University of California, Los Angeles Expected Graduation June 2027

B.S. Statistics and Data Science, Minor in Mathematics GPA: 3.5, Major GPA: 3.8 *Los Angeles, CA*

- **Relevant Coursework:** Mathematical Statistics, Probability Theory, Statistical Computing, Statistical Optimization, Abstract Linear Algebra, Differential Equations, Vector Calculus, Data Analysis and Regression, Experimental Design, Game Theory

Experience

Researcher, UCLA Department of Mathematics Spring 2026

Researcher *Los Angeles, CA*

- Researching statistical arbitrage in cryptocurrency markets: building survivorship-bias-free dynamic universe across 100+ liquid spot pairs, constructing normalized spread series, and testing for mean reversion via cointegration analysis with walk-forward evaluation.

Lutz Consulting Group Nov 2024 – Present

Founder *Las Vegas, NV*

- Analyzed 138 months of client's transaction data; built Gradient Boosting forecasting model with engineered lag, rolling, and seasonal features, achieving 11.5% MAPE, 58% improvement over naive baseline on 24-month holdout validation, uncovering $550K+ in dead stock for liquidation.

New York Life Insurance Company June 2025 – Aug 2025

Data Analyst Intern *New York, NY*

- Built RAG-enabled underwriting chatbot prototype using company underwriting manual as knowledge base; designed controlled experiments measuring LLM sycophancy, producing the team's first empirical model behavior finding that drove production prompt engineering policy. Finalist in company-wide AI Innovation Competition.

Projects

tinynccl - from-scratch RDMA-based PyTorch backend | *C++, CUDA, libibverbs, PyTorch* 2026

- Architected and shipped a from-scratch C++ GPU collective-communication library built directly on RDMA primitives (libibverbs), with pluggable TCP and verbs transports and an all-reduce kernel proven correct under 1000-iteration stress runs across multi-machine, multi-GPU configurations.
- Drove end-to-end PyTorch integration by authoring a custom c10d backend subclass, enabling DistributedDataParallel pipelines to route gradient synchronization transparently through the library; validated by training a 10.7M-parameter transformer across multiple GPUs with bit-exact cross-rank weight consistency.

Quantitative Research Backtester & Platform | *Python, AWS, TypeScript, C++, FastAPI, PostgreSQL* 2025 – Present

- Built data pipeline spanning equity, crypto, macro, and fundamental data across 13,000+ instruments, feeding a co-developed C++ backtesting engine with custom DSL and per-fold walk-forward validation; full-stack platform exposes strategy IDE, real-time charting, and MCP AI integration over FastAPI microservices.
- Architected autonomous research pipeline with model-tiered routing across 27 MCP-orchestrated tools: cheap, fast models pre-screen and terminate weak strategies, larger models drive structured reasoning and validation; designed L0/L1/L2 context hierarchy with prompt caching that cut per-run cost from ˜$10 to ˜$0.50.
- Secured seven-figure AUM allocation from institutional partner (Mazer Group); live pilot generating positive alpha against S&P 500 benchmark.

Skills

Languages: Python, C++, TypeScript, Go, SQL, JavaScript, R, C

Systems & Infrastructure: AWS, PostgreSQL, Docker, FastAPI, Linux, Git, CUDA, RDMA / libibverbs, distributed systems, microservices, REST APIs, MCP server development, Redis, SQLite, Supabase

ML & AI: PyTorch, RAG, LangChain, Hugging Face, FAISS, Pandas, NumPy, Scikit-Learn, prompt caching, multi-agent orchestration, Claude Code

Leadership: Eagle Scout | President, Bruin AI Research Club (Led DataFest team, individual problem-solving award 2023) | Lambda Chi Alpha Executive Board