Rob Schleusner

U.S. Citizen • raschleusner@gmail.com • linkedin.com/in/robschleusner

About

Engineer who pairs systems depth with data-driven analysis, turning large-scale technical measurement into decisions. Energized by latency-sensitive, data-intensive problems — instrumenting and analyzing network and market-data systems, running controlled experiments, and surfacing patterns in large time-series datasets that others miss.

Languages: Advanced – Python, Java, SQL; Intermediate – Rust, C++, C, JavaScript, Scala, HTML/CSS
Tools: Pandas, NumPy, SciPy, DuckDB, Jupyter, AWS, Redshift, SQLite, Redis, Docker, Git, PyTorch, TensorFlow
Concepts: Market microstructure, time-series analysis, probability, networking, optimization, quantitative analysis

Professional Experience

Amazon Inc. **Austin, Texas**
Software Engineer II / Supply and Resource Planning 10/2025 to Present
Software Engineer I / Supply and Resource Planning 02/2024 to 10/2025
Built optimization systems for large-scale driver procurement and scheduling, translating research into production planning models and tooling across North America and Europe.

- Led development of optimization system used by cross-functional planning teams, reducing runtime by 99% and enabling $105M in annual savings across NA/EU operations.
- Partnered with operations and business stakeholders to launch a scalable onboarding framework, increasing planned weekly volume by 20%.
- Designed and built new volume-aware cost-modeling for contracts saving $19M annually.

Dominant Strategies **Austin, Texas**
Software Engineer 08/2023 to 02/2024
Worked on the core protocol team for the sharded blockchain Quai. Built network observability tooling and drove a messaging protocol migration that significantly reduced peer-to-peer latency.

- Built Quai-Stats, a network observability platform that automated collection and visualization of live network metrics across roughly 300 nodes — throughput, block and consensus timing, and peer health.
- Reimplemented peer-to-peer messaging on the libp2p stack, measuring and reducing end-to-end network latency while improving message-delivery reliability across the sharded network.

Neuraflash **Atlanta, Georgia**
Software Engineering Intern / AI & Einstein Solutions Team 05/2021 to 08/2021
Led development and assisted in the design of Getty Images' iStock Einstein Case and Sales Chatbot. Designed and built an internal website that teams use to create large JSON packages. Worked on call to maintain and troubleshoot live chatbots.

- Increased customer satisfaction by 30% on average through chatbot maintenance.
- Automated JSON package generation, saving the sales team roughly 2 hours per week.

Education

Georgia Institute of Technology 05/2025
Master of Science in Computer Science GPA: 4.00
Specialization: Machine Learning

Georgia Institute of Technology **Atlanta, Georgia**
Bachelor of Science in Computer Science, Highest Honors 05/2023
Specializations: Artificial Intelligence and People (UI) GPA: 3.71
Minor: Computing & Business through the Denning Technology & Management Program

Projects

Prediction Market Execution & Analytics Platform 02/2026 to Present
Independent research and trading effort on prediction markets (Kalshi, Polymarket) — full stack from ingestion and signal research through live execution and monitoring.

- Built a Rust execution engine (pre-signed auth, warm HTTP/2 pool, lock-free pre-armed orders) firing on Kalshi and Polymarket in parallel, 35 ms goal-to-venue-ack; beat the market reprice on the majority of goals across 3 live matches.
- Live-deployed an in-venue event-driven strategy on soccer prediction markets: 9 goals, 25k contracts, +$1.4k net on $\leq$$1k deployed; every match followed by a venue-verified post-mortem that drove risk-control fixes.
- Ingested and queried 130M+ Kalshi trades with DuckDB to study market calibration; surfaced systematic mispricings by category, volume, and time-to-close.

Cryptocurrency Trading & Research 06/2021 to 04/2022
Discretionary crypto trading driven by self-directed research — arbitrage, NFT valuation modeling, and on-chain analysis.

- Compounded a personal account from $2K to $400K (200x) over 10 months in which BTC finished roughly flat.
- Built a valuation pipeline for NFT racehorses: scraped 35K+ marketplace sales and full race/odds histories into Postgres, then modeled expected net earnings by odds and bloodline to find horses priced below their future cash flows.

Multilingual: Spanish (conversational), and English (native)