

Arthur Zakrzewski

773-600-0150 | arthurzakrzewski2028@gmail.com | linkedin.com/in/arthur-zakrzewski | github.com/arttomzak

EDUCATION

Northwestern University

Expected June 2028

B.A. in Computer Science and Mathematics

Evanston, IL

- **GPA:** 3.71/4.0 | **SAT:** 1550
- **Relevant Coursework:** Data Structures and Algorithms, Operating Systems, Computer Systems, Kernel Software Development, Design and Analysis of Algorithms, Machine Learning, Proofs, Graph Theory, Probability

EXPERIENCE

Capital One

June 2026 – Present

Software Engineer Intern

Chicago, IL

- Developed an enterprise data-quality framework for 15,000+ datasets from 25+ sources, mitigating partner risks
- Built a Databricks pipeline automating execution of 100+ rules on unregistered data, persisting results to OneLake
- Engineered a Python utility with OAuth to trigger batch job executions via API, eliminating manual workflows
- Refactored 5 legacy notebooks to Python packages deployed via CI/CD, enabling 200+ automated daily job runs

Prescience Lab

March 2026 – Present

Undergraduate Research Assistant

Evanston, IL

- Co-authored research paper on CHERI-based memory safety for kernel-bypass I/O, submitted to ROSS @ SC26
- Built memory-safe kernel-bypass framework for a Solarflare NIC, cutting round-trip latency 43% over an ioctl path
- Wrote an ef_vi-style userspace driver that posts descriptors, rings MMIO doorbells, and busy-polls the event queue
- Designed a thin kernel stub handling interrupt-driven control events and passing CHERI capabilities to userspace

Student Holdings

June 2025 – June 2026

Software Engineer Intern

Evanston, IL

- Led development of a coupon platform, building a React Native app with secure authentication and payment flow
- Architected a maintainable Postgres database schema, streamlining coupon updates to reduce developer overhead
- Developed a FastAPI backend to process webhooks and manage user data, securely isolating admin functions
- Integrated RevenueCat for subscription handling, syncing purchase events with Supabase as a source of truth

ElectronX

August 2025 – September 2025

Data Science Intern

Chicago, IL

- Built a DuckDB-backed Streamlit dashboard querying millions of price records to visualize market volatility trends
- Developed a scalable GeoPandas pipeline parsing Shapefiles to map 2,000+ counties to electricity market regions
- Automated SEC 10-K parsing with BeautifulSoup and Pandas, mapping 500+ electricity utilities to parent entities
- Engineered a fuzzy string matching utility to match companies' generation portfolio projects to 200+ legal entities

PROJECTS

Low-Latency Orderbook and Matching Engine | C++

- Built a limit orderbook and matching engine handling 23M ops/sec at 241ns p99 latency on a day of market data
- Designed a 0-allocation hot path with a fixed-capacity object pool and intrusive linked lists for O(1) cancellation

Counter-Strike Case Profitability Simulator | C++, Python, SteamWebAPI

- Built a multithreaded Monte Carlo simulator to estimate long run Counter-Strike case opening profitability
- Integrated Python scripts to retrieve live skin prices from third-party APIs and visualize expected returns

ProductiviCLI | Rust, PostgreSQL, Neon

- Built a cross-platform CLI productivity tracking tool with Neon-hosted PostgreSQL session data persistence
- Developed a TUI featuring a live session stopwatch and daily/weekly insights queried from historical data

TECHNICAL SKILLS

Languages: C, C++, Rust, Python, TypeScript, JavaScript, SQL

Technologies: Linux, FreeBSD, FastAPI, PostgreSQL, React, React Native, Node, Spark, Pandas, NumPy

Developer Tools: Neovim, Git, GitHub + Actions, Docker, GDB, Valgrind, CMake, AWS, Databricks, Supabase