

Alon Florentin

Machine Learning Engineer — Robotics

Tel Aviv, IL — relocating; open to all US locations • abf386@nyu.edu • (310) 977-8883

github.com/Alonerism • linkedin.com/in/alon-florentin-226103237 • U.S. Citizen — no sponsorship required

Summary

Machine learning engineer (NYU M.S. Data Science, RL coursework) building robotics-adjacent systems end-to-end — ARKit/LiDAR spatial computing on iOS, an RL trading policy validated in backtesting, vehicle-routing/fleet-optimization services, and multi-agent LLM pipelines operated 24/7 in production with cost and reliability telemetry. Owns the full loop from model to deployed, monitored product.

Education

New York University, Center for Data Science

New York, NY

M.S. in Data Science

May 2025

- Coursework: reinforcement learning, natural language processing, applied machine learning for finance, big data

New York University, CAS / Courant Institute

New York, NY

B.A./M.A. in Behavioral Science (joint degree, with honors)

May 2023

- Honors: Maximal Grant, Dean's Undergraduate Research Fund Grant

Experience

Founder

Tel Aviv / Remote

Nika Tech Solutions (AI/automation studio)

June 2025 – present

- Built iOS spatial-computing apps (Swift/ARKit) — a LiDAR room-scanning app extending Apple's RoomPlan framework with scan persistence, human detection, and unit/UI test suites (public repo), and an AR building-model viewer with per-layer visibility and transparency controls.
- Shipped vehicle-routing and fleet-optimization services — wrapped the open-source VROOM VRP solver and built a FastAPI dispatch service with /optimize and per-truck route endpoints plus a route-advance state machine; the same task-allocation problem underlying multi-robot fleet orchestration.
- Designed and operate a multi-agent LLM deal-analysis pipeline (PropHunt) in production 24/7 — 4–10 coordinated AI agents deep-analyze each shortlisted property; deployed on DigitalOcean with per-call LLM cost telemetry, an API-error watchdog, and 36/36 unit tests gating every deploy.
- Engineered server-side event pipelines for a production CRM (Next.js, Postgres) — Conversions-API passback with event-ID deduplication and verified delivery, webhook-driven SMS/email automation (Twilio), HMAC-signed session auth, and nightly automated database backups.
- Built a fully automated media pipeline — JSON spec → generative image/TTS synthesis → programmatic Remotion (React) render → scheduled multi-platform publishing at ~\$0.35 and ~1.5 minutes per finished video, iterated against audience-retention data.

Teaching Fellow

New York, NY

NYU Center for Data Science

Sept 2024 – May 2025

- Courses taught: Big Data (MS), Introduction to Data Science (MS), Principles of Data Science (BS).
- Designed and deployed a model predicting student grades and weak spots from twice-weekly qualitative feedback using NLP (spaCy), XGBoost, and SHAP.

Senior Data Science Intern

New York, NY

Uma Capital (private equity asset manager)

May 2024 – Sept 2024

- Adapted my NYU RL trading research into an internal strategy-research tool — ARIMA and LSTM forecasting with an actor-critic trade policy, reaching 145% annualized return (profit factor 2.1) in backtesting.
- Built a Monte Carlo engine to simulate default risk and IRR distributions for the \$299M Reno–Tahoe ConRAC project; analysis secured unanimous approval for a \$180M term loan.
- Built a Snowflake data warehouse integrating FAA flight data, BTS passenger traffic, and airport IoT sensors; automated pipelines cut data wrangling by 40 hours/quarter and enabled real-time IRR tracking.

Research Assistant (PI Dr. Pascal Wallisch)

New York, NY

NYU FoxLab

Sept 2021 – May 2024

- Led data processing and statistical modeling across behavioral-finance and psychopathology studies; managed 100 GB+ of behavioral data with PySpark, SQL, and distributed compute.
- Co-authored two research manuscripts (PsyArXiv preprints: psychopathy-screening validation; affluence and prosocial behavior) and contributed statistical modeling across additional lab studies, applying Bayesian inference and psychometric frameworks to high-dimensional behavioral data.

Junior Data Science Intern

Los Angeles, CA

Archway Equities (real estate investment)

May 2023 – Sept 2023

- Built LightGBM-powered Monte Carlo simulations integrating wildfire, flood, and sea-level-rise hazard data (CalFIRE, FEMA, CMIP-6) into DCF models with PD/LGD across \$700M+ in coastal real estate; reduced projected financial risk by 27% and supported underwriting for a \$74M acquisition.

Management Intern to the Founder & CEO

New York, NY

Roomrs (real estate investment)

May 2022 – Sept 2022

- Engineered a PySpark data lake and default-risk model for NYC multifamily acquisitions; powered \$40M in private-credit term sheets via Monte Carlo rent-freeze simulations.
- Built an XGBoost 90-day vacancy model (AUC 0.86); reduced average unit downtime by 12 days and increased NOI by 4%.

Skills

Programming: Python, SQL, TypeScript/JavaScript, Swift, C++, Bash, Git

ML & RL: PyTorch, reinforcement learning (PPO, A2C, DQN), XGBoost, LightGBM, PySpark, spaCy, Ray, MLflow, Bayesian inference, Monte Carlo simulation

Spatial & 3D: ARKit, RoomPlan, Unity, Blender, Three.js

LLM & Agents: Anthropic/OpenAI APIs, multi-agent orchestration, MCP server development, evaluation & cost telemetry

Data & Infrastructure: Airflow, AWS, Snowflake, Postgres, FastAPI, production deployment & monitoring (DigitalOcean, Vercel, cron/CI)

Web & APIs: React, Next.js, Node.js, REST & webhook systems

Languages: English (native), Hebrew (fluent), Spanish (conversational)