

Adrian Erlikhman

(213) 330-6884 · erlikhman.adrian@gmail.com · adrianerlikhman.is-a.dev · [LinkedIn](#) · [GitHub](#)

PROFESSIONAL EXPERIENCE

Firstness – Software Engineering & ML Intern · [\[firstness.ai\]](http://firstness.ai) · reporting directly to the founder 2026 – Present

- Work directly under the founder of Firstness, an AI research engine that grows a research question into a map of branching hypotheses, writes each answer from the peer-reviewed literature with a citation on every sentence, and checks a new hypothesis against published prior work.
- Refining how the platform connects into the research stack: the literature databases and reference managers researchers already work in, so an answer arrives inside the tools they use rather than in one more tab.
- Building the demos that walk a new user through what the engine does, why every claim carries a citation, and where each answer can be checked against its source.

Kiddom – Machine Learning Internship & Mentorship · [\[kiddom.co\]](http://kiddom.co) · with Director of Data Science Flora Xu 2026 – Present

- Designing machine-learning and NLP approaches to teacher-facing problems in K–12 assessment: which problems are worth modeling at all, what student data a school can realistically share, and what result would count as evidence rather than a demo.
- Built **Why Wrong** against one of them. It reads a class's answer sheets, names the misconception behind each wrong answer (naive Bayes over answer patterns, then clustering on the residual errors), and groups the class by shared mistake, so a teacher sees *why* an answer was wrong and which students to reteach together. F1 0.87 on planted misconceptions. [\[live\]](#) [\[code\]](#)
- Designing its classroom evaluation with an LAUSD teacher: teacher review of the item bank and the misconception labels, a pilot on de-identified responses from a single class, then a written comparison of where the tool and the teacher agree, where they do not, and what I would change.

Fjor Venture Capital – Investment Intern · seed and pre-seed AI fund 2024 – Present

- Ran diligence on 20+ early-stage AI startups and wrote 10+ full investment memoranda – market sizing, competitive positioning, technical architecture and founder evaluation – for partner-level decisions.
- Led the screen of Y Combinator's Summer 2025 batch with Managing Director Tatiana Janko: sourcing, first-pass screening and the competitive-landscape work behind each shortlisted company.
- Drafted the fund's monthly LP updates and investor-facing materials.

Alliance for SoCal Innovation – Venture Outreach & Pipeline Intern · nonprofit · SoCal startup ecosystem 2024 – 2026

- Built the 250+ firm investor database in HubSpot – VCs, angel syndicates and accelerators – and defined the schema it still sorts on: organization category, funding stage, location and point of contact.
- Led the backend redesign of the venture pipeline with the Chief Innovation Officer, and rewrote the founder intake form so inbound companies route to partners including Expert Dojo and the Long Beach Accelerator.
- Maintained the life-sciences segment the review process runs on, tracking rounds as they closed, and sat in on the SVP Life Science pilot committee's pitch reviews.

VENTURES & CIVIC TECHNOLOGY

AIML-LI (AI/ML Literacy Initiative) – Co-Founder & Lead · [\[aiml-initiative.org\]](http://aiml-initiative.org) · grew out of the LACES AIML Club 2023 – Present

- Co-wrote a free AI-literacy course any teacher can run in twelve 30-minute sessions on the Chromebooks a class already has, plus a 12-week version in six units: students train and break real models, study bias in cases that actually happened, and ship an AI project of their own.
- Wrote the material a non-specialist teacher needs to run it cold – Colab notebooks, instructor guides, student packets and assessments – all free and open-source, so no school pays a license and no lesson depends on a lab. Taught in LACES AP Computer Science, spring 2026.
- Took the course to LAUSD's Ed Tech & Innovation office (Sept 2026), which is carrying it to the district's Chief Academic and Chief Information Officers toward a standing AI-literacy elective. Advised by Philo Juang (Google DeepMind) and Amir Zohrenejad (Data Herald, YC W21); mentored by aiEDU.
- Organizing the LA Student AI Summit & Hackathon for 200+ LA public-school students: keynotes from Google DeepMind, Microsoft, Google and Snap, with Snap, Jane Street, MongoDB, Confluent, Discovery Education and Hacker Fund behind it, under The Hack Foundation as fiscal sponsor.

SafeJew – Co-Founder · [\[safejew.org\]](http://safejew.org) · JFEDLA Teen Innovation Grant 2024 – Present

- Shipped a live map of antisemitic incidents across Greater LA, in production at safejew.org: 800 incidents, 100 of them individually sourced with a citable link on every pin, plus 594 on 33 university campuses.
- Wrote a dependency-free LA gazetteer that geocodes the places named in news prose, so a daily GitHub Action keeps the map current without a human in the loop. Next.js, TypeScript and Supabase; rebuilt as V2 for a Hillel International pitch.
- Draws on ADL H.E.A.T. Map records, news reporting and community submissions, with an anonymous reporting pipeline behind it; distributed through the Jewish Federation of Los Angeles' Community Security Initiative, which funded the build.

- **3rd place internationally, Code for Transportation** (Young Coders’ Sphere). Scores 431,599 LA street blocks for street-crime risk from 85,634 filtered LAPD incidents, then runs A* over them entirely in the browser, with no server, across the 668 LAUSD schools.
- Found that the afternoon walk home is a student’s riskiest hour, and that riding the bus cuts a route’s modeled exposure by up to 80% while getting there faster. Deciding which crimes count was the call the whole thing turned on: keeping only incidents in public space dropped about 81,000 indoor records that would have condemned ordinary residential blocks. [\[live\]](#) [\[code\]](#)

RESEARCH

Whose Line Is It Anyway? Auditing Self- and Other-Recognition in LLM Judges

In submission, TACL

Co-first author, with R. Erlikhman, M. Tarekgn; advised by Philo Juang (Google DeepMind)

- Introduced a peer baseline that scores each LLM judge’s self-attribution rate against how often the other judges name it on the same text, separating self-recognition from text that any model can identify, a confound prior self-recognition studies leave in.
- Nine judges, 1,270 texts, over 19,000 judgments: the usual test flags 15 of 37 cases, 8 survive a false-alarm control, 3 beat the pooled peers and none beats its strongest single peer (Claude names itself on 33 of its 38 essays; GPT names Claude on 32).
- An 18-feature interpretable stylometric classifier recovers the true author 86.3% of the time, above every judge, so the signal is in the text and most judges simply do not use it. Reproducible in one command, with no API key.

Do LLM Fingerprints Survive Outside English?

Accepted, IEEE BigData 2026 HS Symposium

Co-first author

- Collected 840 essays from five frontier models across seven languages, English to Hindi, and attributed them with 21 interpretable features computed from Universal Dependencies parses: 59–72% accuracy in every language against a 20% chance floor, with no decline detected as a language’s resources fall.
- Character n-grams are more accurate within a language, but the interpretable features transfer between languages better (49% against 41%), and attribution still holds at 57–70% on machine-translated text.
- Across 4,195 judge calls, four of the five models attribute authorship at chance (20–25%) while Claude’s self-recognition holds even in Hindi. [\[abstract\]](#) [\[poster\]](#) [\[code\]](#)

Curated to Win: Certifying Reward-Model and Judge Selection Against Evaluation-Set Curation

Accepted, ARRSI Workshop @ AACL-IJCNLP 2026

Sole author

- Proposed a curation certificate for choosing a reward model or LLM judge by validation accuracy: how much of the set could be removed before the winner changes. In benchmark weight it is at least the top-two gap; counted in items it is exact.
- RewardBench’s leader survives the removal of any 11 of 2,985 items, and removing 12 particular items makes the runner-up the winner; linear programs show the bound is tight near the top of RewardBench and RewardBench 2. On 200-item RewardBench samples, picks with a top-two gap above 2 points match the full-benchmark winner 87% of the time (43% otherwise). [\[pdf\]](#) [\[code\]](#)

How Robust Are the Legatum Prosperity Index Rankings?

A Monte Carlo and Sensitivity Analysis of Weighting, Normalization, and Aggregation

Accepted with revisions, Journal of High School Science

Co-first author, with Ryan Erlikhman

- Reproduced every published score for 167 countries exactly, then varied the choices the index hides: 10,000 draws under each of three weight priors, and a 131,072-run Sobol’ design over weighting, normalization and aggregation.
- Derived rank leverage, $c^2(1 - \rho^2)$ – the only part of a weight’s influence that can actually reorder countries – and showed it tracks Sobol’ total effects (Spearman 0.91–0.99) at a fraction of the compute. The ranking holds as a whole (median Spearman ≥ 0.9899) but not mid-table, where the median 90% rank interval spans 36 places. [\[pdf\]](#)

Does Portfolio Optimization Beat Splitting Evenly? Markowitz vs. the 1/N Portfolio

Accepted with revisions, Oxford Journal of Student Scholarship

With Ryan Erlikhman

- Equal weights beat Markowitz in 74% of 250 calibrated simulations (Sharpe 0.70 against 0.51), about 3 points of return a year at equal risk. The loss traces to estimating expected returns, and shrinkage closes it only with roughly two decades of stationary data – more history than most investors have. [\[pdf\]](#)

Advanced Mathematics Research

· induction, number theory, combinatorics and game theory, with PhD candidate Kirill

Kovalenko

2022 – Present

COMPETITIONS & BUILDS

Ledger, VISION HACK: South LA – 1st place, Council District 9 Champion, U18 · Hacker Fund × City of LA Vision Lab

Sep 2026

- Led the twelve-hour build across the vision pipeline, rule engine and frontend, and wrote the two-pass vision pipeline, the dashboard and renewal calendar, and the static fallback. Ledger puts the 18 programs a corner store answers to, across three levels of government, on one renewal calendar, and scores a photo of a wholesale order record against the SNAP stocking rule.
- Built to undercount: anything the pipeline reads below 0.75 confidence is held back and listed for the owner rather than counted, so a blurry photo shortens the count instead of padding it. English and Spanish throughout. [\[live\]](#) [\[code\]](#)

eDNAtlas, Decode the Ocean Hackathon – 1st place · Lovable × United Nations 2026

- Built and shipped the winning entry: a live map that turns environmental DNA sampling into a plain-language health score for coastal sites, with the raw evidence and Darwin Core downloads behind every score. [live] [code]

Citadel High School Terminal – Finalist, top 8 of 45 teams · Citadel | Citadel Securities, hosted by Correlation One 2026

- Wrote “dunerscore,” a Python tower-defense strategy that reads the board each turn and prices every wall, turret and attack lane, then adapts to the opponent’s last layout; shipped a new version after every round of replay post-mortems, on a three-person team. [about]

Y Combinator Startup School – Admitted · sub-5% acceptance · \$30,000+ in partner credits 2026

- Worked the company-building program end to end: validation, user discovery and go-to-market. [about]

More builds

- **Quant:** HMM regime-aware portfolio optimizer, LSTM equity forecaster benchmarked against ARIMA and linear baselines, FinBERT market-sentiment signal, and imbalance-aware fraud detection scored on PR-AUC (0.36 against a 0.28 baseline) with a recall-weighted operating threshold.

TALKS & PRESENTATIONS

NCSS 2026, National Consortium of STEM Schools – Invited Speaker, Student Panel on AI · San Diego, CA Dec 2026

- Invited by aiEDU onto the student panel in its pre-conference session “Transform Your School’s AI Readiness”: how students actually experience AI, and what the roughly 100 educators in the room need to know.

LAUSD Innovation Expo – Presenter · the AIML-LI curriculum, alongside Google, Microsoft and Canva

LAUSD Student Innovation Challenge – Invited Presenter · a \$10,000 district challenge

EDUCATION

Los Angeles Center for Enriched Studies (LACES) – LAUSD magnet school · Los Angeles, CA Class of 2027

- GPA 5.2 weighted / 4.0 unweighted · SAT 1520 (Math 770, EBRW 750) · 2× AP Scholar with Distinction
- 17 AP courses, 11 completed and 6 this year, across computer science, mathematics and the sciences.
- Dual enrollment at L.A. Southwest College: Linear Algebra (MATH 270) · Programming in Python (CS 119) · Python Programming for Data Science and Machine Learning (CS 121, planned)

LEADERSHIP & ATHLETICS

Competitive Fencing, Épée – USA Fencing A26, Team USA · LA International Fencing Center

- U.S. No. 49 Junior (U20) and No. 78 Senior; peak No. 17 Cadet (U17); No. 1 in Region 4. Represented Team USA at the 2026 Cadet Pan-American Championships in Bogotá (6th). In active recruitment with MIT.
- Top 32 Division I, National Championships · Top 8 Cadet, SJCC · Top 16 Cadet, October North American Cup · 4× Junior regional champion · 3× All-American and 3× All-Academic First Team
- Coach the club’s younger épéeists, including my brother, through their first competitive season.

Jewish Student Union – President · 50+ members 2023 – Present

- Run weekly programming with NCSY and organized a school-wide Holocaust Remembrance assembly featuring a survivor, with StandWithUs.

STEMsters – Co-Founder & President · 30+ volunteers · 4 elementary schools 2022 – Present

- Place LACES volunteers in four elementary schools to teach hands-on coding, robotics and engineering to kids who would not otherwise get it.

HONORS

- National Merit Semifinalist (about 16,000 of 1.3M+ entrants) · JFEDLA Teen Innovation Grant · SIFMA Stock Market Game, top 5% statewide as a solo entrant · MTAC piano, Level 9 candidate, five state-honors exams

SKILLS

Code: Python, TypeScript/JavaScript, Java, SQL, LaTeX; Next.js, React, Supabase/PostgreSQL, GitHub Actions, Git

ML & statistics: PyTorch, scikit-learn, TensorFlow/Keras, pandas, NumPy, SHAP, Stanza, FinBERT; LLM evaluation, stylometry, Monte Carlo and Sobol’ sensitivity analysis, hypothesis testing, time series, HMMs, A* search

Certifications: Certiport IT Specialist in Software Development, Cybersecurity, Python, Data Analytics, AI and HTML/CSS; Deep Learning Specialization (Andrew Ng, Coursera); NVIDIA Deep Learning Fundamentals; Kaggle × Google Generative AI Intensive

Languages: English, Russian (fluent), Spanish (intermediate)