

alexwritesprograms@
gmail.com
813.220.1066

Alex "AC" Chojnacki (he/him)

[LinkedIn](#)
[GitLab](#)
[Blog](#)
[Google Scholar](#)

Staff-level ML/platform engineer specializing in production ML, applied AI, healthcare NLP, data science platforms, and geospatial analytics. I take real-world data products from prototype to production, and I have a lot of fun doing it.

Summary

- » 8+ years building production ML, applied AI, geospatial analytics, and data products across healthcare, utilities, climate resilience, and public infrastructure.
- » Led ML platform, founding engineering, and CTO work spanning MLOps, technical strategy, hiring, vendor management, customer pilots, and multi-tenant SaaS delivery.
- » Published applied ML research at KDD and related workshops, including award-winning work on ML-guided lead service line replacement in Flint, Michigan.

Work Experience

Data Science Platform Engineer Paradigm Health **May 2023 — Present**

- » Served as platform lead for a 15–25 person data science and machine learning organization, owning Databricks enablement, GitLab CI, Terraform infrastructure, Datadog observability, and PHI-enabled Kubernetes services.
- » Built the v1 clinical trial matching workflow, powering 5–10 million inferences over one year to match cancer patients with clinical trials using structured criteria configured by informatics users.
- » Bootstrapped and maintained AI services for human-in-the-loop study conduct workflows, combining OCR, retrieval, document automapping, and LLM extraction to pre-populate ECRF fields with source citations.
- » Established production ML engineering standards across Python applications: repo templates, package management, CI/CD, IaC, testing, runbooks, alerting, monitoring, and deployment workflows.
- » Led AI engineering enablement through mentorship, design docs, ADRs, vendor evaluation, and an AI guild focused on RAG, medical LLMs, AI tooling, and software development best practices.

Co-founder, CTO, Rhizome Data **October 2022 — May 2023**

- » Co-founded grid resilience analytics company applying causal modeling, machine learning, and geospatial data to distribution utility investment planning.
- » Delivered Vermont Electric Co-op pilot modeling outage consequences and five-year load forecasts for risk pricing and regulatory rate-case strategy.
- » Developed technical strategy for resilience analysis, fundraising, hiring, product documentation, and geospatial prototype design.

Founding Engineer, BlueConduit **September 2020 — October 2022**

- » Served as founding engineer and engineering lead for my tenure, productionizing ML research from the Flint water crisis for lead service line identification and remediation planning.
- » Built command-line ML workflows for training, evaluation, tuning, batch prediction, audit logging, and retraining of service line material classifiers.
- » Developed geospatial prioritization workflows that aggregated block-level lead risk, helping utilities target inspections and reduce costly false-positive excavations.
- » Led delivery of a multi-tenant SaaS product for 20–30 municipal water customers using Next.js, ECS, Aurora Postgres, AWS CDK, Cognito, Python, pandas, and scikit-learn.
- » Built and managed the engineering function across DevOps, frontend, full-stack, contractors, and Google.org fellows, including equitable hiring, vendor management, and a nationwide lead service line map.

Contractor Developer, self-employed **May 2021 — August 2021**

- » Built data ingestion and normalization workflows for NYC real estate datasets, including PLUTO and BoBA, to support parcel-level geospatial analytics.
- » Built full-stack prototype features using Clojure/ClojureScript, a Datalog graph database, tileserver, and map-based search frontend.

Machine Learning Engineer / Software Engineer, Urbint **January 2018 — July 2020**

- » Prototyped and maintained a utility damage-prevention ML product that contributed roughly half of company ARR in 2019.
- » Built shared Python ML tooling for feature engineering, model evaluation, batch scoring, and production deployment across the data science team.
- » Prototyped and deployed computer vision models for a customer-facing video classification product.
- » Built full-stack product features in React, Haskell, and a custom graph database-backed application.

Developer, University of Michigan & the City of Flint **March 2017 — January 2018**

- » Built production web application used by City of Flint and Michigan National Guard teams to record service line inspections, coordinate a dozen replacement crews, and maintain the material inventory source of truth across 55,000 parcels.
- » Developed active learning workflows and lead service line prediction models using field labels from 35,000 occupied parcels to guide data-driven inspection and replacement prioritization.
- » Modeled lead risk with XGBoost, logistic regression, and hierarchical Bayesian geospatial smoothing, including block-level prioritization for inspection and remediation planning.
- » Adapted data collection, reporting, and field workflows to operational feedback; this approach later formed the basis for BlueConduit deployments in other cities.

Earlier Experience **2015 — 2016**

- » Software engineering intern at Nordstrom, modeling POS fraud incidence with the payments team; received a return offer.
- » Web developer at University of Michigan CAEN, independently designing and implementing a RESTful JSON API for the UM Data Warehouse.

Technical Expertise

ML/AI	Production ML, MLOps, model evaluation, model serving, RAG, agents, NLP, computer vision, time series, causal inference
Data & Geospatial	SQL, Postgres/PostGIS, Databricks, Spark, Parquet, DuckDB, geospatial analytics, tileservers, Mapbox/MapLibre
Infrastructure	AWS, GCP, Kubernetes, Terraform, Docker, CI/CD, serverless, monitoring, alerting, Datadog
Languages	Python, SQL, TypeScript/JavaScript, Clojure/ClojureScript, Haskell, Julia
Developer Tools	git, bash, emacs, Nix, package management, testing, code review, technical design docs

Ask me about data science platforms, ML/AI model serving, RAG systems, Postgres and PostGIS, geospatial data products, workflow orchestration, vector stores, DuckDB, Kubernetes for data science, BI tooling, and why developer experience matters for AI adoption. I've got opinions and would love to compare notes.

Education

- University of Michigan, Ann Arbor, MI** **September 2013** — **June 2017**
- » LS&A, B.S. in Computer Science
 - » Machine Learning, Web Applications, Cognitive Science, French

Research & Teaching Experience

- Research Assistant, University of Michigan** **2015** — **2017**
- » Researched active learning methods for Flint lead service line replacement, including Python simulations comparing proposed inspection policies to historical performance.
 - » Built experimental web instrumentation to study native advertising efficacy.
- Instructional Aide, EECS 398/201, Computing for Computer Scientists** **2016** — **2017**
- » Lectured on `git`, Linux servers, debuggers, `bash`, developer tooling, and practical software engineering workflows.
 - » Helped craft the original syllabus for an experimental course that became EECS 201, Essentials of Computer Science; personally taught 600+ students over 3 semesters.
- Grader & Tutor, EECS 183, Introduction to Programming Concepts** **2014** — **2015**
- » Assisted students in one-on-one meetings and graded homework and projects for correctness, design, and code style.

Selected Publications

- » J. Abernethy, **A. Chojnacki**, et al. “**The Search for Lead Pipes in Flint, Michigan.**” *KDD Conference 2018. Awarded Best Student Paper in Applied Data Science Track.*
- » J. Abernethy, **A. Chojnacki**, et al. “**A Data Science Approach to Understanding Residential Water Contamination in Flint.**” *KDD Conference 2017.*
- » J. Stroud, **A. Chojnacki**, J. Abernethy. “**The Michigan Data Science Team: A Student Organization for Machine Learning Challenges.**” *NIPS 2016 Workshop, Challenges in Machine Learning: Gaming and Education.*
- » J. Abernethy, C. Anderson, **A. Chojnacki**, et al. “**Data Science in Service of Performing Arts: Applying Machine Learning to Predicting Audience Preferences.**” *Bloomberg Data For Good Exchange 2016.*

Campus Leadership

- Michigan Data Science & ML Organizations, University of Michigan** **2015** — **2017**
- » Held leadership and organizer roles with Michigan Data Science Team, Michigan Machine Learning Reading Group, and Michigan Student Artificial Intelligence Lab.
 - » Grew Michigan Data Science Team by 40 active members and 400 newsletter subscribers while organizing prediction challenges with external sponsors and research collaborators.