

STRENGTHS

-  **Trust and safety judgment**
During policy reviews, translated uncertain signals into testable labels and safer thresholds; partners became a go-to sounding board.
-  **Evidence-led evaluation**
Built error reviews around human labels and calibration checks; one careful comparison prevented an unsupported production launch.
-  **Cross-functional clarity**
Explained model tradeoffs in plain language for policy and product partners, earning confidence during staged deployment decisions.
-  **Practical GenAI design**
Tested retrieval, citations, prompts, and guardrails in a reviewer workflow; useful limits mattered as much as useful output.
-  **Reliable delivery habits**
Versioned datasets, documented experiments, and monitored drift across releases; teammates relied on clean handoffs when pressure rose.

SKILLS

<u>Machine learning</u>	<u>Python</u>	<u>PyTorch</u>
<u>TensorFlow</u>	<u>SQL</u>	<u>NLP</u>
<u>Model evaluation</u>	<u>GenAI workflows</u>	
<u>LLM orchestration</u>		
<u>Data exploration</u>	<u>ML pipelines</u>	
<u>Cross-functional work</u>		
<u>Model monitoring</u>	<u>Responsible AI</u>	
<u>Technical documentation</u>		

SUMMARY

Associate Machine Learning Engineer with nearly four years applying Python, PyTorch, TensorFlow, SQL, NLP, and **model evaluation** methods across trust and safety, content quality, and data products. Builds reproducible **ML pipelines**, monitoring dashboards, and **human-in-the-loop** review tools that help policy specialists identify risky activity consistently. Explores LLM and GenAI workflows using retrieval, structured prompts, guardrails, and multi-step orchestration, with close attention to transparency, auditability, and responsible AI. Converts experiments into clear documentation for engineering, **data science**, policy, and product partners. Brings practical judgment, steady communication, and genuine interest in making digital communities safer and more inclusive.

EXPERIENCE

Machine Learning Engineer

SignalHarbor Labs  January 2025 - Present  Columbus, Ohio

- Owns trust and safety ML initiatives spanning text classification, ranking, monitoring, GenAI review support, and audit-ready delivery. Partners with policy, product, data science, and engineering groups across production workflows.
- Built **PyTorch** classifiers for spam and harassment signals across 18 million weekly records.
 - Created **Python** and **SQL** feature pipelines with validation, versioning, and reproducible training workflows.
 - Developed monitoring views for precision, recall, drift, review volume, and threshold changes.
 - Prototyped retrieval-augmented **GenAI** support with citations, prompt controls, and restricted-action guardrails.
 - Compared PyTorch and **TensorFlow** experiments against human labels for safer operating thresholds.
 - Documented **model cards**, experiment decisions, limitations, and audit artifacts for cross-functional handoffs

Machine Learning Engineer

CivicPattern Analytics  July 2023 - December 2024  Raleigh, North Carolina

- Delivered **NLP** and data science solutions for digital community services, guiding **model evaluation**, scheduled scoring, compliance reporting, and staged deployment decisions. Worked closely with policy analysts and technical stakeholders.
- Implemented Python NLP pipelines that routed high-risk user reports into specialized review queues.
 - Consolidated event, text, and reviewer data with SQL, pandas, provenance, and privacy controls.
 - Evaluated **TensorFlow** and **PyTorch** baselines using confusion matrices, calibration, and stratified validation.
 - Automated scheduled scoring and performance reporting for weekly quality and **compliance** reviews.
 - Translated ambiguous **policy** guidance into labeled examples, escalation rules, and acceptance criteria.
 - Presented residual risks and model tradeoffs, supporting a staged deployment decision

Data Science Analyst

BlueOak Data Cooperative  September 2022 - June 2023  Durham, North Carolina

- Supported supervised learning studies through data preparation, exploratory NLP analysis, baseline modeling, reviewer tooling, and responsible-use assessment. Maintained clear records that helped senior scientists move from inquiry toward reliable evaluation.
- Prepared customer-support datasets with SQL and **Python** for supervised learning studies.
 - Built exploratory notebooks that surfaced abuse-report language patterns for NLP feature design.

LANGUAGES

English	Native
Spanish	Intermediate

MY CAREER



- Machine Learning Engineer at SignalHarbor Labs (1.7 Years)
- Machine Learning Engineer at CivicPattern Analytics (1.4 Years)
- Data Science Analyst at BlueOak Data Cooperative (9 Months)

- Compared logistic regression, tree models, and TensorFlow text models with **precision, recall**, and F1.
- Created a **human-in-the-loop** labeling interface for rationale and disagreement capture.
- Tested prompt-based LLM classification against rules, documenting bias and reproducibility concerns.
- Maintained data dictionaries, experiment notes, and deployment-readiness checklists for **auditability**

PROJECTS

Policy Evidence Copilot 📅 2025

A useful safety tool does more than produce an answer. It shows its evidence. Built a retrieval-augmented workflow that assembled policy passages, cited sources, applied guardrails, and routed uncertain cases for human review.

Risk Signal Evaluation Lab 📅 2024

Better moderation starts with better questions. Created a reproducible evaluation environment for NLP classifiers, threshold studies, calibration checks, confusion analysis, and reviewer-capacity scenarios.

LEADERSHIP & AWARDS

- SignalHarbor Labs Responsible AI Recognition, 2025; honored for practical audit documentation and safer GenAI review workflows.
- CivicPattern Analytics Delivery Award, 2024; recognized for turning ambiguous policy guidance into usable evaluation criteria.
- North Carolina State University Computer Science Leadership Award, 2022; celebrated peer mentoring and applied ML work.

EDUCATION

Bachelor's Degree in Computer Science

North Carolina State University 🎓 GPA: 3.7 📅 2022 📍 Raleigh, North Carolina

Coursework: Machine Learning, Natural Language Processing, Database Systems, Software Engineering, Data Mining

CERTIFICATIONS

- Google Advanced Data Analytics Professional Certificate 📅 2023
- AWS Academy Machine Learning Foundations 📅 2024
- Responsible AI Practices, Microsoft Learn 📅 2024

TECHNICAL SKILLS

- Programming Languages:** Python, SQL, Bash
- Machine Learning Frameworks:** PyTorch, TensorFlow, scikit-learn
- Data Analysis Tools:** pandas, Jupyter, NumPy
- NLP Methods:** Text classification, ranking, tokenization
- GenAI Systems:** LLMs, retrieval-augmented generation, prompt evaluation
- ML Operations:** ML pipelines, model deployment, scheduled scoring
- Model Monitoring:** Precision, recall, class drift
- Data Engineering:** Feature engineering, dataset validation, data versioning
- Responsible AI:** Guardrails, bias review, human-in-the-loop
- Documentation Standards:** Model cards, audit artifacts, data dictionaries
- Cloud And Containers:** AWS, Docker, AWS Machine Learning Foundations
- Development Tools:** Git, REST APIs, modular components

PROFESSIONAL AFFILIATIONS

- Association for Computing Machinery member; contributes thoughtful ML practice and peer exchange.

- Partnership on AI community participant; supports responsible AI, transparency, and inclusive technology discussions.

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST