

Machine Learning Engineer

We are seeking an experienced Machine Learning Engineer to help us build state-of-the-art AI tools that help people make sense of complex legislative and regulatory documents. You will play an important role in harnessing the power of LLMs and Machine Learning to help policy professionals in nonprofits, governments, and enterprises do their work more accurately and effectively. Plural Engineering is a diverse, remote-first team that values autonomy, collaboration, and continuous learning. We are looking for a driven and curious individual that will be able to own their part of the AI work, which will require significant research and innovation.

This is a remote position that requires collaboration with teammates in the U.S. Eastern time zone. This contract position has the opportunity to grow into a permanent position based on both individual and company performance. The contract will be for a minimum of three months, with options to extend contract terms or convert to full-time after a successful initial engagement. This position reports to the Head of Engineering.

Job Overview

We are looking for an experienced **Machine Learning Engineer**. This engineer will be responsible for the implementation and maintenance of Python-based AI/ML tools in Plural's product portfolio. Working with our Data Science Lead, you will:

- Do research and development on AI tools for policy professionals
- Work with policy experts to build evaluation frameworks for AI projects
- Expand Plural's existing suite of AI tools to new data types and use cases
- Collaborate with the Web team to define contracts for web-based UIs that control AI tools and agents
- Document the design and result of experiments efficiently and clearly
- Stay up-to-date on AI research and make recommendations to the department on which new techniques to adopt
- Quickly iterate on demos and proofs-of-concept for enterprise policy applications
- Grow demos & experiments into flexible, reliable applications suitable for production use

Required Skills

- Previous experience building Machine Learning models or statistical analysis

- Basic knowledge around NLP, Deep Learning and Large Language Models.
- Fluency in Python and SQL
- Experience with python libraries like Numpy, Pandas, Sklearn and Pytorch
- Can “TAP in” to Plural’s core values:
 - Embrace **Transparency**
 - Think **Audaciously**
 - Be **People-First**

Our ideal candidate has

- Experience building and leading ML projects in the wild with multiple stakeholders and unclear outcomes
- Experience training and deploying Deep Learning models, including working with CUDA and Pytorch
- Experience working with Open Source NLP and Large Language Models, including libraries like Transformers, Ollama, etc.
- Experience building features with proprietary language models from OpenAI, Anthropic or Google.
- Familiarity with Agentic AI techniques such as: tool use, custom workflows, few shot learning, personas, REACT, human-in-the-loop strategies, etc.
- Experience deploying ML models as services with software like Docker, FastAPI, Flask.
- Experience with Data Engineering services like Airflow, Databricks, etc.
- Experience with Cloud Architecture like GCP, AWS and Azure
- Experience with ML evaluation suites and metrics
- The desire to work in a fast-paced R&D environment, iterating from idea to proof-of-concept to working tools
- Experience in the policy or legal tech domain

Location

- This is a fully remote position — Required working hours are based on U.S. Eastern time zone

Compensation

- Starting at \$6k per month, negotiable based on experience and location

Commitment to Equal Opportunity

Plural is an equal opportunity employer committed to fostering a diverse, inclusive, and equitable work environment. Individuals seeking employment at Plural are considered without regard to race, ethnicity, religion, national origin, age, gender, gender identity, gender expression, sexual orientation, marital status, medical condition, ancestry, military or veteran status, physical or mental abilities, social-economic background, political affiliation, thinking styles, life experiences, and any other characteristic protected by applicable law.

About Us

[Plural](#) is a venture-backed startup building technology for the future of civic engagement and political transparency. Our flagship software product makes it easy for leading organizations to find and analyze legislation and regulations around the world. We are committed to growing the digital commons of open civic data through our [Plural Open project \(github\)](#). Some of our AI features include AI Bill Summarization, Bill Classification by Policy Topics and Compliance Monitoring Agents.

We are a remote-first company, with team members in multiple US locations.

Plural's Product Offerings:

- Plural's **Bill Tracking and Intelligence** combines premium policy data, real-time updates, and powerful AI features to give teams a competitive edge. Users can access state and federal legislation, hearing and lawmaker data, and collaborate with other users on bill analysis and reporting.
- Launched in 2025, the **Policy Analyzer** is for researchers and lobbyists who need answers on the go. This cutting-edge AI-powered tool converts long, dense bill text into digestible summaries – making legislative research and advocacy easier.
- Still in development, Plural's **Regulatory Lifecycle Management** solution helps corporate teams understand, respond to, and comply with changes to the regulatory landscape. Our enterprise-grade AI automates and streamlines complex processes, saving businesses time and money by analyzing public policy data specific to the enterprise, identifying risks, and automating compliance workflows.
- With Plural's **Policy Data API**, anyone can access and automate our open-source collection of legislative data. Our API offers flexibility and power to small teams and enterprise organizations alike, whether for analyzing trends, tracking legislation, or building custom reports. The API is free to access at low volume.