Associate AI Engineer - Mission Innovation Lab

HEJ higheredjobs.com/search/details.cfm

Type: Full-Time
Posted: 17 days ago
Category: Software Engineer/Programmer
Military

Carnegie Mellon University

At the SEI AI Division, we conduct research in applied artificial intelligence and the engineering questions related to the practical design and implementation of AI technologies and systems. We currently lead a community-wide movement to mature the discipline of AI Engineering for Defense and National Security.

As our government customers adopt AI and machine learning to provide leap-ahead mission capabilities, we

- build real-world, mission-scale AI capabilities through solving practical engineering problems
- discover and define the processes, practices, and tools to support operationalizing AI for robust, secure, scalable, and human-centered mission capabilities
- prepare our customers to be ready for the unique challenges of adopting, deploying, using, and maintaining AI capabilities
- identify and investigate emerging AI and AI-adjacent technologies that are rapidly transforming the technology landscape

Are you creative, curious, energetic, collaborative, technology-focused, and hard-working? Are you interested in making a difference by bringing innovation to government organizations and beyond? Apply to join our team.

Overview

As an associate AI Engineer who thrives at the intersection of deep-learning research and production-grade software development, you will translate cutting-edge AI concepts into robust, mission-scale solutions for the warfighting community. You will work comfortably with large-scale foundation models such as GPT and LLaMA, designing and deploying agentic workflows, as well as apply and advance traditional ML research and engineering across domains such as natural language processing, computer vision, time series forecasting, and other predictive analytics. You will collaborate closely with senior researchers, software engineers, and government sponsors to define problem statements, iterate on experimental designs, and deliver secure, reliable AI capabilities that meet stringent mission requirements.

The Mission Innovation Lab within the SEI's AI Division works with the defense and national security community to translate the "recently possible" in AI into reliable mission and warfighting capabilities.

Key Responsibilities

Design, develop, and fine-tune a variety of AI models.

Design autonomous agents and multi-step pipelines using LangChain, ReAct, tool-calling, or custom orchestration; employ the Model Context protocol to manage stateful interactions.

Build Retrieval-Augmented Generation pipelines that combine external knowledge bases with LLMs to improve factual accuracy for warfighting applications.

Implement end-to-end data pipelines, ETL processes, and back-end services (Python, C/C++, Java) that feed data to models.

Create CI/CD pipelines for model training, validation, containerized deployment (Docker/Kubernetes), and security scanning; maintain model registries, monitoring, and version control of context protocols.

Produce rapid prototypes, run benchmarks, and conduct robustness/adversarial testing in realistic environments.

Work closely with senior ML engineers, software developers, and government customers; mentor junior staff and contribute to design reviews and documentation.

Stay current with emerging LLM architectures, agentic paradigms, PEFT/LoRA methods, and AI-safety techniques; translate new research into operational capabilities.

Required Qualifications

Bachelor's degree in Computer Science, Machine Learning, Statistics, Applied Mathematics, or a related field with at least 3 years of relevant experience (or an M.S. with 1 year)

- You will be subject to a background investigation and must be able to obtain and maintain an active Department of War (DoW) security clearance.
- You must be able and willing to work onsite 5 days per week at an SEI office in either Pittsburgh, PA or Arlington, VA.

Proficiency in Python and at least one compiled language (C/C++ or Java); experience with REST/GraphQL APIs and containerization.

Strong grasp of ML theory (supervised, unsupervised, reinforcement learning) and evaluation metrics.

Hands-on experience fine-tuning LLMs and using frameworks such as Hugging Face Transformers, LangChain, or comparable agent tools.

Familiarity with building RAG pipelines (vector stores, dense/sparse retrievers).

Experience applying PEFT/LoRA methods (e.g., LoRA, adapters) to large models.

Understanding of Model Context protocols for managing model state across multi-turn interactions.

Experience building evaluation frameworks, benchmarks, or data quality pipelines

Experience with TensorFlow, PyTorch, or JAX; knowledge of data-pipeline tools (Airflow, Prefect, Ray) is a plus.

Awareness of DevSecOps practices (CI/CD, GitOps, container security scanning, model-registry concepts) is desirable.

Desired Experience

Deploying LLM APIs (FastAPI, gRPC) at scale, handling latency and load balancing.

Building multi-tool agents, planner-executor loops, or tool-calling pipelines for complex decision-making.

Conducting adversarial testing, implementing input sanitization, and contributing to AI-safety research.

Utilizing GPU/TPU resources, mixed-precision training, and distributed training frameworks such as DeepSpeed or ZeRO.

Prior work on defense, intelligence, or government-focused AI projects and familiarity with DoW acquisition or compliance processes.

Contributing to open-source AI and ML libraries, agentic frameworks, or context-protocol implementations.

Knowledge, Skills, & Abilities

Analytical thinking: decompose complex AI problems into tractable components and iterate rapidly.

Strong written and verbal communication skills for documenting designs and presenting results to technical and non-technical stakeholders.

Proven teamwork: collaborate in interdisciplinary groups, mentor peers, and contribute to shared codebases.

High curiosity and autonomy: proactively explore emerging technologies and integrate them into mission work.

Location

Pittsburgh, PA

Job Function

Software/Applications Development/Engineering

Position Type

Staff - Regular

Full time/Part time

Full time

Pay Basis

Salary**More Information:**

- Please visit **"Why Carnegie Mellon"** to learn more about becoming part of an institution inspiring innovations that change the world.
- Click here to view a listing of employee benefits
- **Carnegie Mellon University is an Equal Opportunity Employer/Disability/Veteran**.
- **Statement of Assurance**

Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state, or local laws or executive orders.