



AI Engineer (CHR/AI-20250812)

USA / Contract / Remote

We're looking for an **AI Engineer** to help design, build, and deploy GenAI solutions that power our platform. This is a hands-on role for an engineer who enjoys working with modern LLM workflows, creating efficient pipelines, and shipping production-ready features quickly.

Key Responsibilities

- Build and deploy Python-based GenAI solutions for analyzing and improving ad content
- Design and implement LLM-based pipelines using tools like LangChain, LlamaIndex, and OpenAI APIs
- Own development of RAG systems, vector search, and prompt workflows
- Optimize for low-cost, high-reliability inference (cloud and open-source LLMs)
- Build robust APIs to integrate AI models into production systems
- Move fast: ship features, not just prototypes

Required Skills & Experience

- 5+ years of experience in Python and AI/ML engineering
- Strong understanding of LLM-based workflows, embeddings, prompt engineering, and RAG
- Experience with a modern GenAI framework: LangChain, LlamaIndex, etc.
- Ability to build and deploy APIs using FastAPI, Flask, or equivalent
- Familiarity with vector databases (e.g. FAISS, Pinecone, Weaviate)

- Solid software engineering habits: clean code, version control, CI/CD
- Experience working with cloud platforms (AWS, GCP, Azure) and LLM cost optimization

Why Join us?

- Work on large-scale AI projects with real-world impact
- High autonomy, high-impact role on a lean, ambitious team
- Shape both the product and technical foundation of AI systems
- Fully remote work with flexibility
- Competitive compensation and opportunities for career growth.
- Continuous learning, experimentation, and exposure to cutting-edge AI tools