

AI Engineer — Solumn AI

Experience: 3–5 years · **Type:** Full-time · **Work week:** 6 days

About Solumn AI

Solumn AI builds the safety and alignment layer for frontier AI systems.

We work with leading AI labs to make advanced models deployment-ready, across the post-training and assurance stack: red teaming, vulnerability discovery, safety evaluations, agent-behavior testing, and the datasets and frameworks that make it measurable. If capability asks "what can the model do?", we ask "what happens when it fails, and how do we prevent it?"

Why this matters: Every technology cycle begins with capability and matures through trust. The internet, cloud, mobile and social all ran the same arc, with security and governance becoming indispensable only after the first wave of scale. AI is crossing that threshold now, and the assurance layer is becoming foundational. We're building the firm that does it to a benchmark standard.

The role

You will train and fine-tune models, and own the engineering behind our platform, benchmarks and scale. This is a hands-on role with engineering-manager scope: you set technical direction, review others' work, and grow a team as we scale.

What you'll do

- Fine-tune and post-train LLMs (SFT, RLHF, RL on verifiable rewards) and measure the results.
- Design and build benchmarks and eval harnesses for agentic and safety behaviour.
- Own the platform that generates, runs and verifies RLE tasks — Dockerized environments, graders, rollout pipelines.
- Scale throughput from hundreds to thousands of tasks a week without losing quality.
- Set engineering standards, review code, and mentor engineers and contributors.
- Work directly with frontier-lab clients on technical requirements and delivery.

What we're looking for

- 3–5 years in ML/AI engineering, with hands-on model training or fine-tuning in production or research.
- Strong Python and PyTorch; working knowledge of Hugging Face, TRL/DeepSpeed/vLLM or similar.
- Experience building evals or benchmarks for LLMs or agents.
- Solid systems skills: Docker, cloud (AWS), CI, distributed jobs.
- Has led projects or people, or is ready to step into engineering-manager scope.
- Strong pedigree: a top-tier institute, or a top AI lab or AI-first company.

Nice to have

- Publications at top ML venues.
- Experience in AI safety, red-teaming or security.
- Built RL environments or tool-calling agent setups.