



AI AND AGENTIC SYSTEMS ENGINEER

Full-time · US (Hybrid)

At Laocoön Health, we are building foundational infrastructure that makes healthcare economics more intelligent, explainable, and humane. This role is for engineers who care deeply about correctness, elegance through simplicity, and building durable, future-proofed systems - modern-day monuments that stand the test of time.

THE ROLE

This role is for engineers who want to build intelligent systems that behave predictably, even under conditions of ambiguity, complexity, and uncertainty. You will design and implement agentic systems that combine structured knowledge extraction, retrieval, and language model-based tooling to support clinical and economic decision-making.

Your job is to build the part of the system that acts. You'll design agents that turn clinical and economic reasoning into real decisions, ones that retrieve the right evidence, weigh it against policy, and submit something a human would sign their own name to.

This goes far beyond automation. The task ahead requires taste and judgment, with decision-making delegated carefully and held to a human-like standard across dynamic, fuzzy clinical states.

The systems you build won't just predict; they'll act and decide, and most importantly, they'll be right when it matters *most*.

WHAT YOU'LL DO

- Develop multi-step, goal-directed clinical agents using frameworks like LangChain, LangGraph, or OpenAI's Agents SDK
- Integrate LLMs with knowledge graphs and structured data; integrate agents into complex, ambiguous clinical workflows
- Design guardrails for determinism, reliability, and explainability
- Build and maintain RAG pipelines, reflection loops, prompt chains, and real-time monitoring, logging, and observability for agent behaviors, errors, and outcomes
- Optimize agent cost-efficiency and latency by managing prompt templates, caching strategies, and model selection (OpenAI, Anthropic, DeepSeek, Gemini, etc.)



- Work with product, data, and engineering teams to integrate agents into existing workflows while maintaining SOC 2 and data governance standards
- Evaluate and improve agent systems against real-world scenarios, edge cases, and usage loads to ensure scalability and accuracy

YOU ARE

- Curious about intelligence *as a system*, not a model
- Skeptical of black boxes, biased to design with human in the loop philosophy
- Motivated by making AI real-world useful, not flashy
- Positive, open- and solution-minded; biased toward action, inclusivity, and radical curiosity

MINIMUM QUALIFICATIONS

- Experience building LLM-powered systems in production
- Strong software engineering fundamentals

PREFERRED QUALIFICATIONS

- Advanced degree (Masters or PhD) strongly preferred (AI, ML, CS, or related field): 2 years of industry experience with a Bachelors can be substituted
- Experience with RAG, agent orchestration, and tool-using models
- Background in healthcare or other high-stakes, zero-trust, failsafe domains