

Adam Rybinski

AI Engineer | LLM & Multi-Agent Systems Specialist

+35679081417 · adam@compose.systems · linkedin.com/in/arybinski · github.com/adamrybinski

TECHNICAL EXPERTISE

Multi-Agent AI & LLM Systems

LangGraph · LangChain · CrewAI · MCP · AutoGPT · ADK (Agent Development Kit) · RAG · Multi-Agent Orchestration · LLM Models (OpenAI, Claude, Gemini, Proprietary) · Prompt Engineering · Few-shot Optimization · Instruction Fine-tuning

AI Platform & Infrastructure

Cloud-agnostic AI/ML solutions using Terraform, Helm, Kubernetes (EKS, AKS, GKE) with GPU-accelerated model hosting, CI/CD/Argo workflow automation, and multi-cloud scale. Vendor-agnostic platforms integrating OpenAI, Claude, Gemini, and custom models for multi-agent orchestration using LangChain, CrewAI, AutoGPT, ADK, and proprietary frameworks with dynamic workload allocation.

Cloud-Agnostic ML Infrastructure & DevOps

Kubernetes · AWS · Azure · On-Prem · Terraform (IaC) · Docker · Argo Workflows · Helm · GPU Node Orchestration · GitOps · CI/CD

ML Ops & Model Serving

GPU-Powered Inference · Model Hosting on K8s · Vector Databases · Pipeline Automation · Model Versioning · Low-Latency Serving · Microservices Architecture

Development & Integration

Python · Node.js · React/TypeScript · REST APIs · GraphQL · TensorFlow · PyTorch · Watson AI · Hugging Face · Real-time Data Processing

PROFESSIONAL EXPERIENCE

AI Engineer & Multi-Agent Systems Specialist, European Patent Organisation

April 2024 – Present

- Designed **multi-agent AI systems** with proven experience with agent frameworks including **LangGraph, LangChain, CrewAI, MCP, AutoGPT, ADK (Agent Development Kit)**, and custom orchestrations for automated patent report generation with specialized agents for research, analysis, and synthesis
- Architected **cloud-agnostic AI platform** on **AWS, Azure**, and **on-premise** using **Terraform IaC, Kubernetes** with **GPU orchestration, Argo Workflows**, and **Helm** for zero vendor lock-in
- Built **end-to-end ML pipelines** with **Docker**, **GitOps** integrating **OpenAI, Claude**, and proprietary models; implemented **GPU-powered serving** on K8s, reducing report generation time by 75%

AI & Cloud Infrastructure Consultant, EarnLab

2023 – 2024

- Architected and deployed **AWS infrastructure** for box company platform using **Kubernetes** consulting and **security-hardened deployments** of **Node.js** services with **Nuxt.js** framework
- Implemented **Vue.js UI** applications with **WebSocket** real-time communication for enterprise logistics platform
- Established **DevOps best practices** with **Terraform IaC, CI/CD pipelines**, and **security compliance** for production deployments on **AWS EKS**

AI Engineer - Agentic Systems, ICS (Intelligent Chat Solutions)

2023 – 2024

- Built **agentic chat system with RAG** utilizing **Cloudflare AI Workers, AI Gateway**, and **open-source LLMs** including **Llama** models for intelligent document retrieval
- Architected **serverless AI infrastructure** on **Cloudflare Workers** with **vector databases** for scalable RAG implementation, serving real-time responses with sub-100ms latency
- Integrated **multi-model orchestration** across **Cloudflare AI Gateway** enabling dynamic model selection between proprietary and open-source LLMs for cost optimization and performance

Freelance AI Engineer & LLM Consultant

April 2022 – April 2024

- Built **multi-persona AI chat** with experience in agent frameworks including **LangChain, AutoGPT, AI Gateway, Claude**, and **OpenAI** for marketing client with specialized customer support personas, serving **10K+ daily users**
- Deployed **cloud-neutral AI solutions** for Fortune 500 with modular architecture across **AWS/Azure**, leveraging **Kubernetes** and **Terraform**
- Implemented **RAG systems** with vector databases for document analysis, utilizing **prompt engineering** and **few-shot optimization**, improving accuracy by 35%
- Developed **React/TypeScript** frontends for AI platforms with real-time data processing and cognitive science-driven human-AI interfaces

Senior Software Engineer, Betsson Group

April 2023 – April 2024

- Architected **RAG-based AI system** using **LangChain** and vector databases for semantic search across codebase documentation for **200+ developers**, reducing search time by 60%
- Built **AI-powered code analysis tools** integrating **OpenAI** with internal repositories; developed **React/TypeScript** interfaces on **AWS** with real-time semantic search
- Implemented vector embedding pipelines for repository indexing with **Python/Node.js** microservices architecture

Senior Software Engineer, Deriv Europe

October 2021 – April 2023

- Developed **React trading interfaces** with real-time data visualization, improving decision-making speed by 25%
- Built **GraphQL/REST APIs** for high-frequency financial platforms; implemented **RAG systems** and **vector databases** for intelligent document search in financial documentation
- Created accessible UIs for complex financial tools used by global trading community with data-driven design

Software Engineer, IBM

January 2017 – 2021

- Developed **AI-powered presentation analysis** using **Watson AI, Speech-to-Text**, and **Image-to-Text** to auto-segment presentations with transcripts and navigation (precursor to YouTube chapters)
- Built enterprise **cognitive AI interfaces** using **Watson** for healthcare and banking, integrating NLP and computer vision APIs on **IBM Cloud** and **on-premise**
- Implemented **multi-modal AI pipelines** combining speech recognition, image analysis, and NLP for automated content understanding; created **React/TypeScript** frontends with real-time transcript generation

EDUCATION

M.S. Cognitive Science

Jagiellonian University, 2017–2019

Thesis: Explainable AI through Deep CNN structure analysis

B.Sc. Neuroscience

Jagiellonian University, 2011–2015

Project: Neural Network Analysis with Graph Theory

LANGUAGES

English (C1) · Polish (Native)