

Agentic AI for Scientists

A 6-week hands-on workshop series

Tired of constantly trying to keep up with new AI tools and wondering what is really happening behind the curtain?

Agentic AI for Scientists is a 6-week hands-on workshop series designed to introduce researchers and engineers to the rapidly evolving landscape of agentic AI. Beginning with the foundations of modern AI, the series traces the progression from classical machine learning and deep learning to large language models and, ultimately, agentic AI systems. Moving from foundational concepts to advanced applications, the workshop explores agent design patterns, multi-agent architectures, post-training, deployment, evaluation, and the emerging frontier of autonomous AI agents for scientific research.

Designed for participants with diverse backgrounds and levels of experience, this series aims to provide not only a strong conceptual understanding of AI agents, but also practical insight into how they can be applied to scientific workflows. Participants will explore how agentic AI can support research automation, data analysis, and knowledge discovery, while gaining hands-on experience through weekly projects and applied exercises. By the end of the series, attendees will be well equipped to understand, build, and evaluate AI agents for real-world scientific and technical use cases. Each session will also feature special guest speakers with domain expertise relevant to the week's topic.

Weekly content

Week 1: GenAI Foundations for Scientists. A brief history of the AI revolution: ML to DL to Transformers to LLMs to AI Agents. A practical introduction to "vibe coding" to leverage coding agents.

Week 2: AI Agents Fundamentals I, Patterns. Core agentic patterns: ReAct, CoT/ToT-style reasoning, reflection and reflexion agents. Tool use, function calling, and information retrieval and grounding with RAG.

Week 3: AI Agents Fundamentals II, Multi-Agent Systems. Designing multi-agent systems for deep research and automated knowledge discovery.

Week 4: Post-Training and Deployment of AI Agents. How to make LLMs "behave" and let them live in the real world.

Week 5: Evaluation and Benchmarking of AI Agents. Is it good enough for you?

Week 6: Towards Full Autonomy. OpenClaw for Scientists.

What you will learn

- How modern AI evolved from machine learning to large language models and AI agents.
- The core design patterns behind agentic systems, including ReAct, reasoning-based approaches, and reflection-style agents.
- How agents use tools, external APIs, and retrieval systems such as RAG to work with real information.
- How multi-agent systems can be designed for research automation, deep research, and knowledge discovery.
- What post-training and deployment involve, including how to improve model behavior and prepare agents for real-world use.
- How AI agents can be evaluated and benchmarked to determine whether they are reliable and useful.
- Emerging directions toward greater autonomy in scientific workflows.

Why it matters

The latest generation of AI systems can do far more than generate text. They can reason, use external tools, retrieve and synthesize information, coordinate with other agents, and support increasingly autonomous workflows. For scientists, researchers, and technical professionals, this creates new opportunities to accelerate discovery, streamline complex tasks, and build intelligent systems tailored to specific domain needs. This workshop series helps participants move beyond basic familiarity with large language models and develop a structured understanding of how agentic AI works in practice.

Instructor

Dr.-Ing. Kerem Delikoyun is a research fellow at TUMCREATE, working at the intersection of biomedical imaging, translational medical technology, and agentic AI. He earned his Ph.D. in Electrical and Computer Engineering from the Technical University of Munich (TUM), where his research centered on AI-driven automated data analysis and biomarker discovery in hematology, with a particular focus on point-of-care diagnostics in acute care. His expertise spans applied machine learning and agentic AI, particularly in medtech applications using LLM-based multi-agent systems. Through this workshop series, he aims to make advanced

AI concepts more accessible, practical, and empowering for scientists who wish to integrate these tools into their own research and work.

Logistics

- **Schedule:** Weekly over 6 sessions.
- **Time:** Wednesday, 2 to 4 PM.
- Participants are encouraged to bring their own laptops for a more interactive learning experience. Any required software, accounts, or preparatory instructions will be shared in advance before the workshop begins.

Materials

Slides, runnable Colab notebooks, and notes for each delivered week are published openly at github.com/krmde/agent-ai-for-scientists-workshop.