

ArtIMis: Artificial Intelligence for Mission

Value Proposition

ArtIMis is an advanced agentic AI ecosystem designed to speed scientific discovery by combining specialized AI agents, scientific foundation models (SciFMs) and rigorous test-and-evaluation workflows into one coordinated platform. It helps researchers move faster from question to insight by automating literature review, hypothesis generation, code creation/execution, simulation support, analysis and reporting across complex scientific problems.

Technology Readiness Level 5-6

IP Information

U.S. Patent pending

Contact Information

licensing@lanl.gov

Overview

ArtIMis operates as a modular system of coordinated agents built around URSA, the Universal Research and Scientific Agent. A user submits a high-level scientific question, and the planning agent breaks it into tasks that can be assigned to specialized agents for research, code generation, simulation support, optimization and analysis. The system can use external tools, scientific data sources and both frontier language models and laboratory-developed scientific foundation models to carry out the work. A built-in evaluation framework then checks outputs against domain-specific scientific standards so the system can refine results rather than generate them. ArtIMis offers a flexible foundation for faster iteration, more consistent workflows and better use of high-value scientific talent for organizations that need scalable research support with stronger validation than an AI assistant can provide.

Advantages

- Speeds up scientific workflows
- Reduces repetitive manual research tasks
- Supports multiple scientific domains
- Combines planning, execution and evaluation in one system
- Works with complex data types and simulation outputs
- Designed for scalable research environments

Technology Description

ArtIMis is a Python-based, agentic workflow environment built on LangGraph and LangChain, with support for command-line use, YAML- or JSON-defined workflows and a web-based interface. The architecture is designed to support stateful, multi-agent execution with controlled task decomposition, feedback loops and structured tool use. That makes the system suitable for scientific workloads that benefit from automation but still require traceability, evaluation, and the ability to combine multiple data types, including simulation outputs, images, text and structured scientific data.

The platform also includes scientific foundation models tailored to domain-specific tasks rather than general text generation alone. Those models are intended to support areas such as partial differential equation surrogates, materials discovery, fracture prediction and subsurface prediction, while the broader workflow layer coordinates the agents that plan, execute, validate and summarize scientific work. The result is an integrated research environment that can be adapted to different scientific domains without rebuilding the workflow from scratch.

Market Applications

- **Materials Science** (alloy discovery, fracture prediction, process optimization)
- **Chemistry and Molecular Design** (chelator discovery, property prediction, screening workflows)
- **Computational Physics** (PDE surrogate modeling, simulation analysis, scientific modeling)
- **Biotechnology and Structural Biology** (protein analysis, binding prediction, scientific data workflows)
- **National Laboratories** (simulation support, model evaluation, research acceleration)
- **Scientific research institutions** (literature review, hypothesis generation, workflow automation)

Learn More

<https://www.lanl.gov/media/news/0416-meet-ursa>