OAK RIDGE National Laboratory

Agentic AI for Nuclear Materials Design: From Tool to Co-Investigator

"When does a workflow become a collaborator?"

Ana Gainaru CCSD
Denise Adorno Lopes FFESD

ORNL Future of Computing Collaboration Catalyst, 22 April 2026

ORNL is managed by UT-Battelle LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Motivation: Why Nuclear Materials?

Why it's hard

Extreme Complexity Across Scales
From atomic defects (picoseconds) to fuel rods (m, years)
Couple physics: radiation, chemistry, mechanics

Fragmented Knowledge
Experiments, simulations, ML models operate in isolation
Integration relies on manual workflows and expert judgement

Slow Qualification Cycle
Conventional sequential testing takes decades to qualify new fuels and architectures

Multiscale Nuclear Materials Behavior

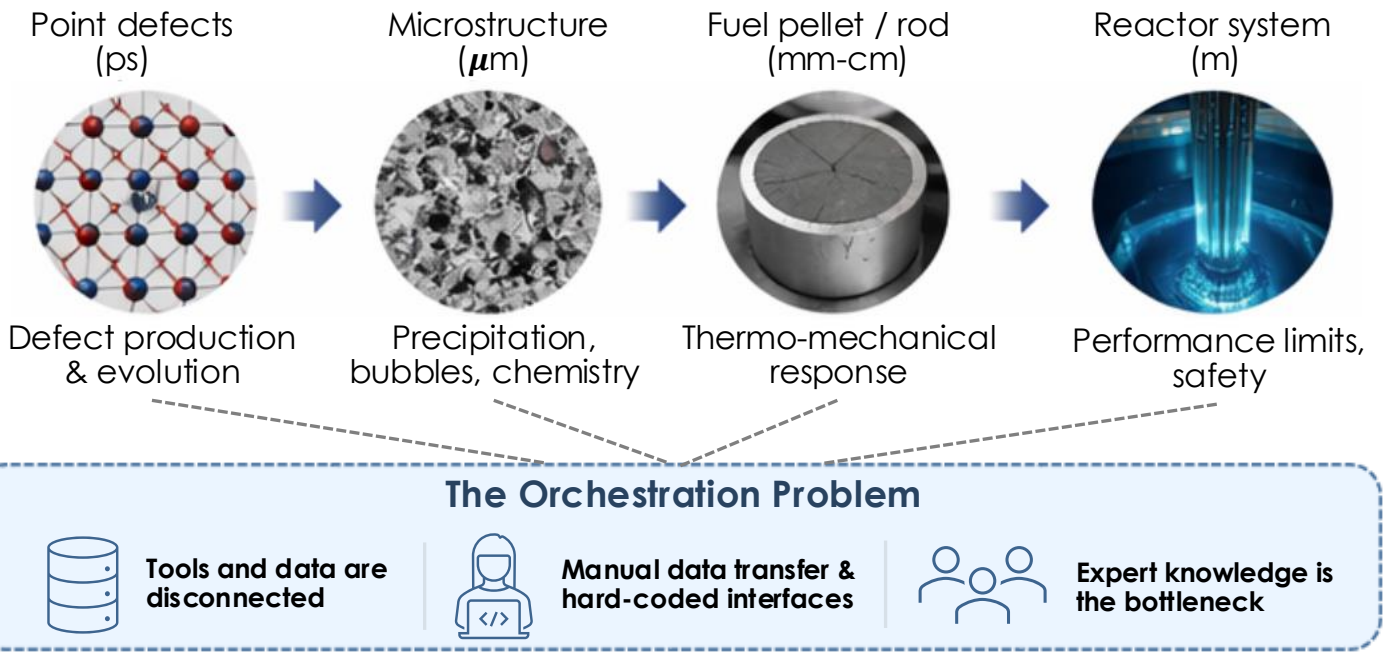

The Orchestration Problem

Tools and data are disconnected

Manual data transfer & hard-coded interfaces

Expert knowledge is the bottleneck

We don't lack models or data-we lack a way to connect them into a coherent scientific workflow

Agentic systems can **dynamically connect models, data and simulations**, and automate hypothesis → simulation → interpretation loops

From months of analysis to minutes

When agents start making scientific decisions,
are they still tools or co-investigators?

What We Mean by "Agentic System"

From fragmented tools to an integrated, human-centered workflow

What Changes with Agentic AI

- **Brings models, data and experiments** together as interoperable modules
- **Automates the scientific workflow** hypothesize – plan – run – analyze – loop
- **Learns and remembers** from prior results, assumptions and failure modes
- **Improves speed and reproducibility** by enabling traceable reasoning

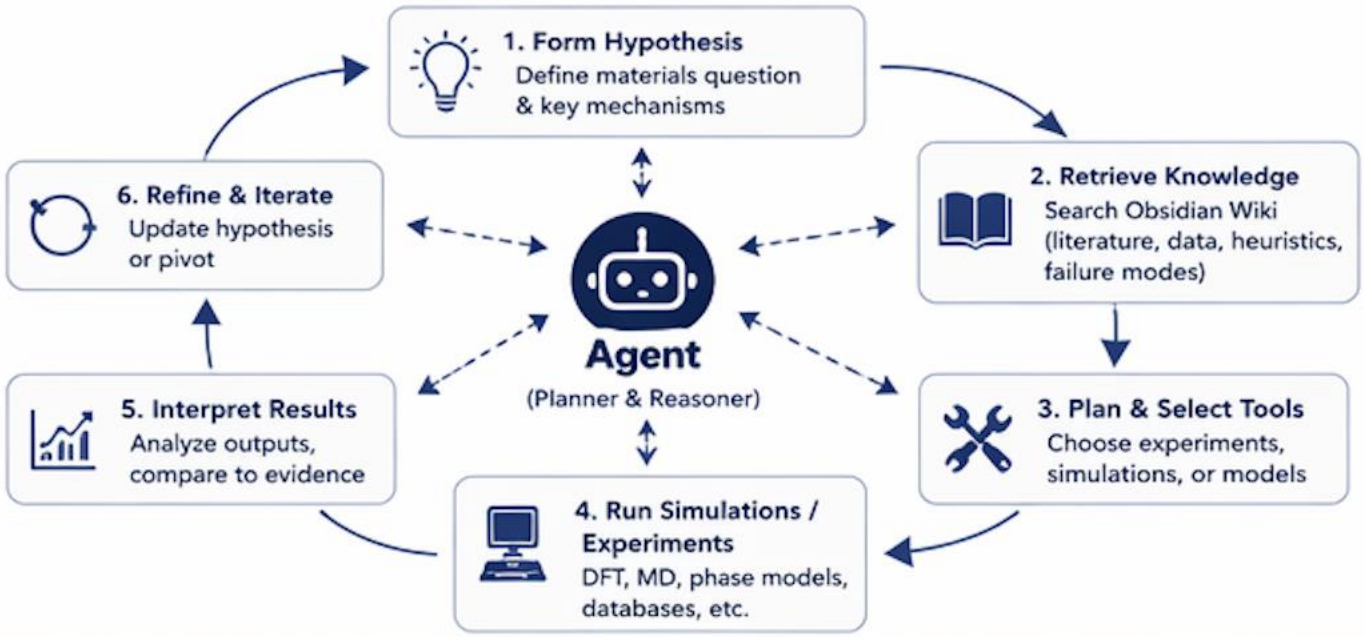

Agentic Framework

LLM-based planner

Knowledge base: Wiki of prior experiments, assumptions, failure modes

Tools: DFT/MD simulation interfaces, materials database, surrogate models

From tools to co-investigation

Agentic AI does not replace scientists, it augments them. The goal is faster, more informed decisions while **preserving human judgment** as the ultimate arbiter of knowledge.

Degrees of freedom

Accuracy/cost tunning through task granularity: precise vs general tasks

Handling the knowledge base: provenance tracking, GNN, recommendation system.

Example Workflow

1.2 GB based on over 500 DFT simulations and 2 million downloaded entries

https://github.com/ORNL-Inria/AERIS

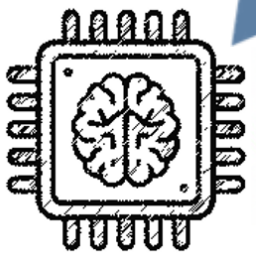

Reproduce the findings of a fission paper

Journal of Nuclear Materials 510 (2018) 331–336

Contents lists available at ScienceDirect

Journal of Nuclear Materials

journal homepage: www.elsevier.com/locate/jnucmat

Stability of U_5Si_4 phase in U-Si system: Crystal structure prediction and phonon properties using first-principles calculations

D.A. Lopes*, V. Kocevski, T.L. Wilson, E.E. Moore, T.M. Besmann

Nuclear Engineering Program, Department of Mechanical Engineering, University of South Carolina, Columbia, SC, 29209, USA

ARTICLE INFO

Article history:
Received 30 May 2018
Received in revised form
6 August 2018
Accepted 13 August 2018
Available online 14 August 2018

ABSTRACT

U-Si systems have recently received considerable attention due to the potential application of U_3Si_2 as a high-density fuel under an accident tolerant fuel initiative. However, the thermodynamic stability of the more recently reported adjacent U_5Si_4 phase is uncertain and could play a significant role in fuel performance. In this work, the enthalpy of formation of the phase predicted by density functional theory (DFT) using the DFT + *U* formalism is used with an evolutionary algorithm (USPEX) to evaluate stability and possible atomic structures for U_5Si_4. The structure of U-Si convex hull phases and the confirmed U_3Si_2 structure were predicted providing confidence in the reliability of the evolutionary algorithm, as well as to obtain a convex hull with comparable enthalpies of formation. Subsequently, the code was 18 and 36 atoms, predicting a 36-atom hexagonal symis the lowest-energy configuration, agreeing with that calculations using the density functional perturbation ed structure is dynamically unstable, exhibiting negative ated that generated shears are directed toward the foralogous to those occupied by carbon atoms in $U_{20}Si_{16}C_3$. phase from assessed U-Si phase equilibria is currently

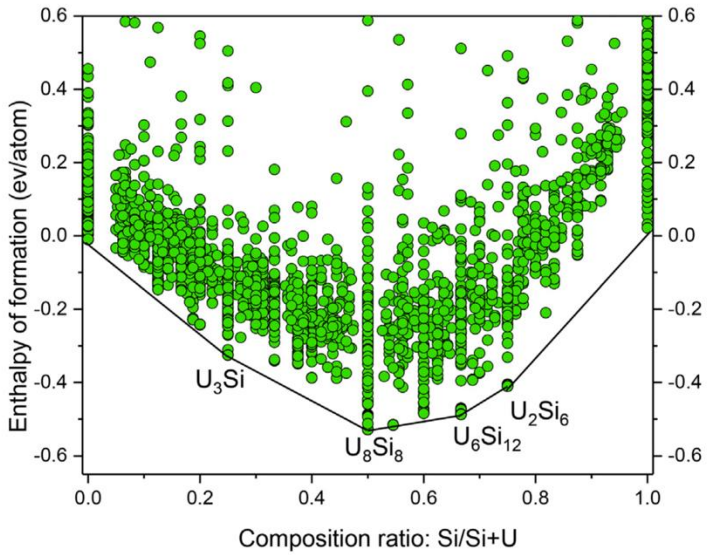

OAK RIDGE National Laboratory

Agentic Framework (AERIS)

AERIS reasoning engine

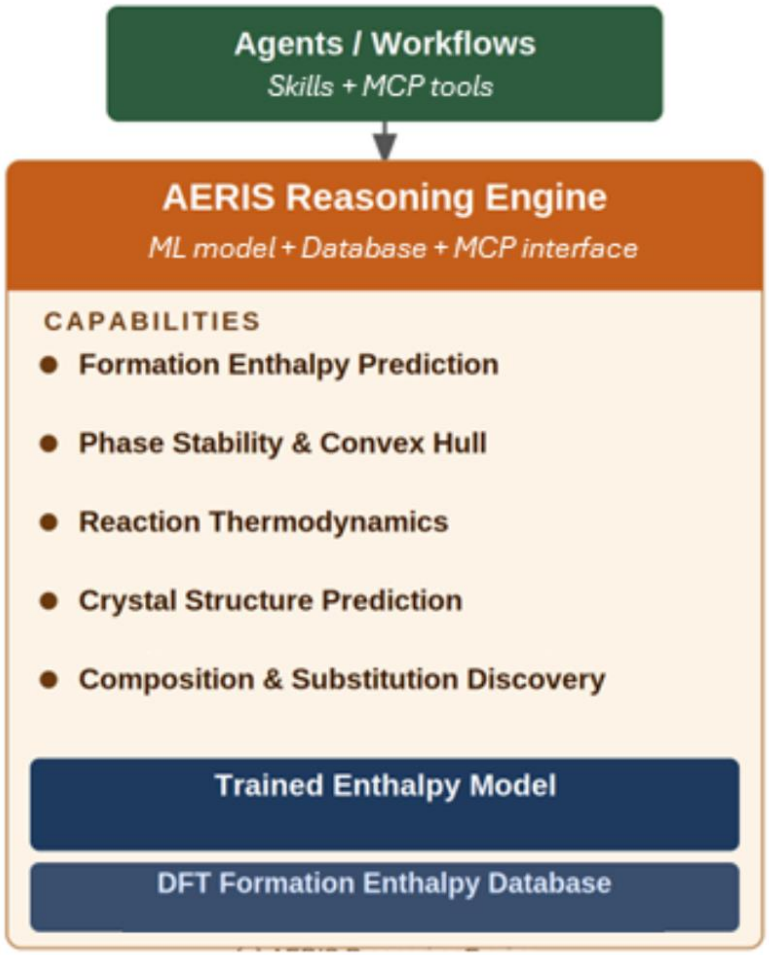

Single model multiple agents

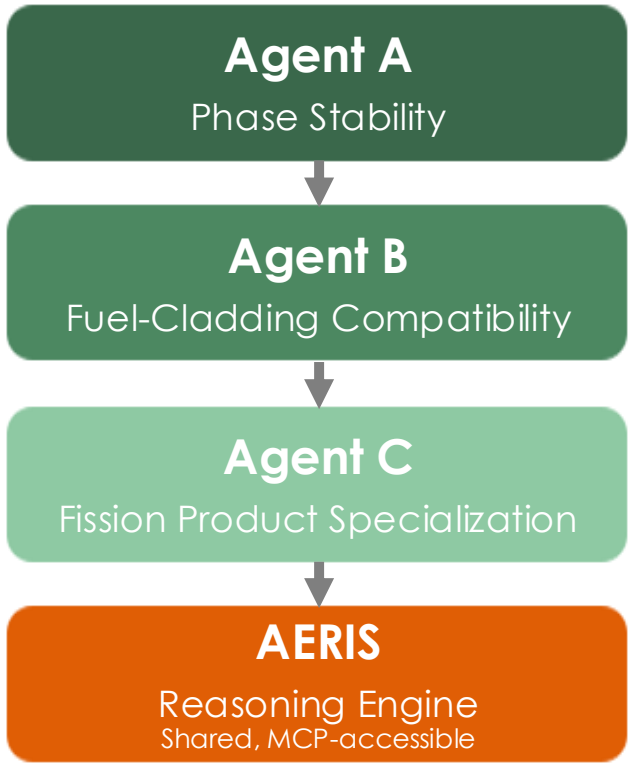

https://github.com/ORNL-Inria/AERIS

Agentic Framework (AERIS)

AERIS reasoning engine

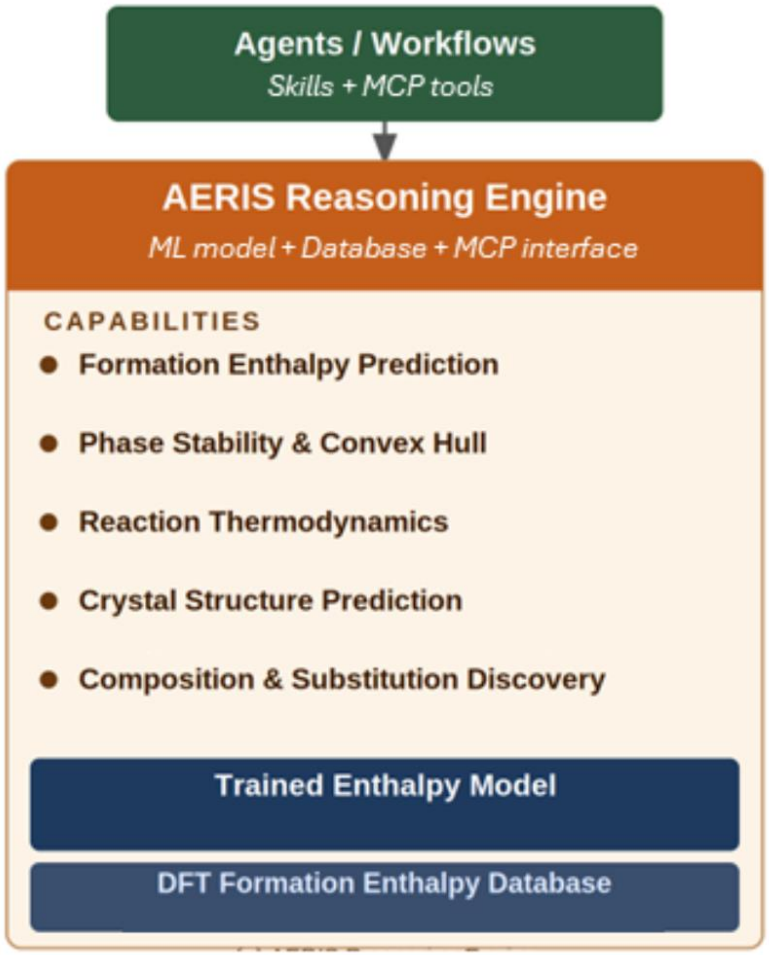

Single agent multiple resources

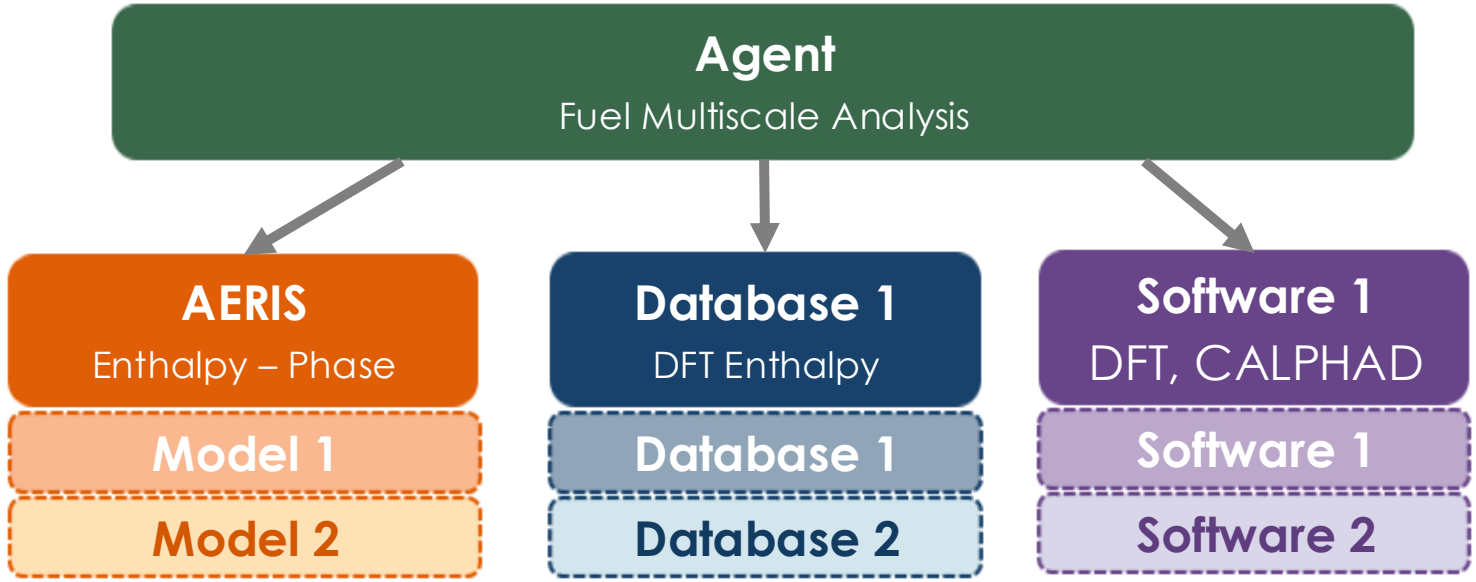

https://github.com/ORNL-Inria/AERIS

Where It Starts Acting Like a Co-PI

Case Study: Reproducing the U-Si Binary Convex Hull

Task given to agent:
"Characterize the U–Si system with formation enthalpies and determine stable phases."

What the agent did autonomously

- Enumerated candidate compositions across U–Si
- Selected relevant structural templates
- Predicted formation enthalpies using AERIS Enthalpy (DFT-based)
- Constructed the convex hull
- Identified stable phases and decomposition neighbors

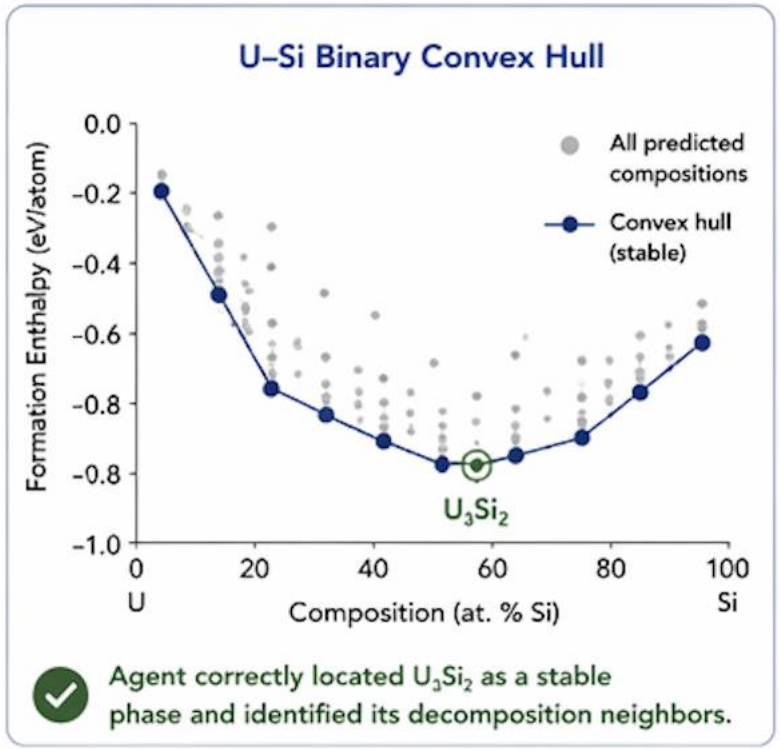

Agent correctly located U_3Si_2 as a stable phase and identified its decomposition neighbors.

Why this is scientific reasoning—not just execution

Chooses which components to evaluate: Defines the search space (hypothesis framing)

Decides which structures are relevant: Encodes physical assumptions (structure selection)

Selects which models to query: Makes methodological choices (model, database)

Determines lowest-energy phases: Evaluates multiple structure prediction and filters based on relevance to the hypothesis

Infers stability & decomposition pathways: Produces scientific conclusions (knowledge generation)

The agent is not just running calculations- it is deciding what counts as evidence and what conclusions to draw

Runtime: minutes | No predefined workflow – the agent determines the sequence from the task | (Adorno Lopes et al., 2025)

Human as Arbiter: Keeping Scientists in the Driver's Seat

Ensuring human judgment remains the ultimate authority on plausibility and knowledge

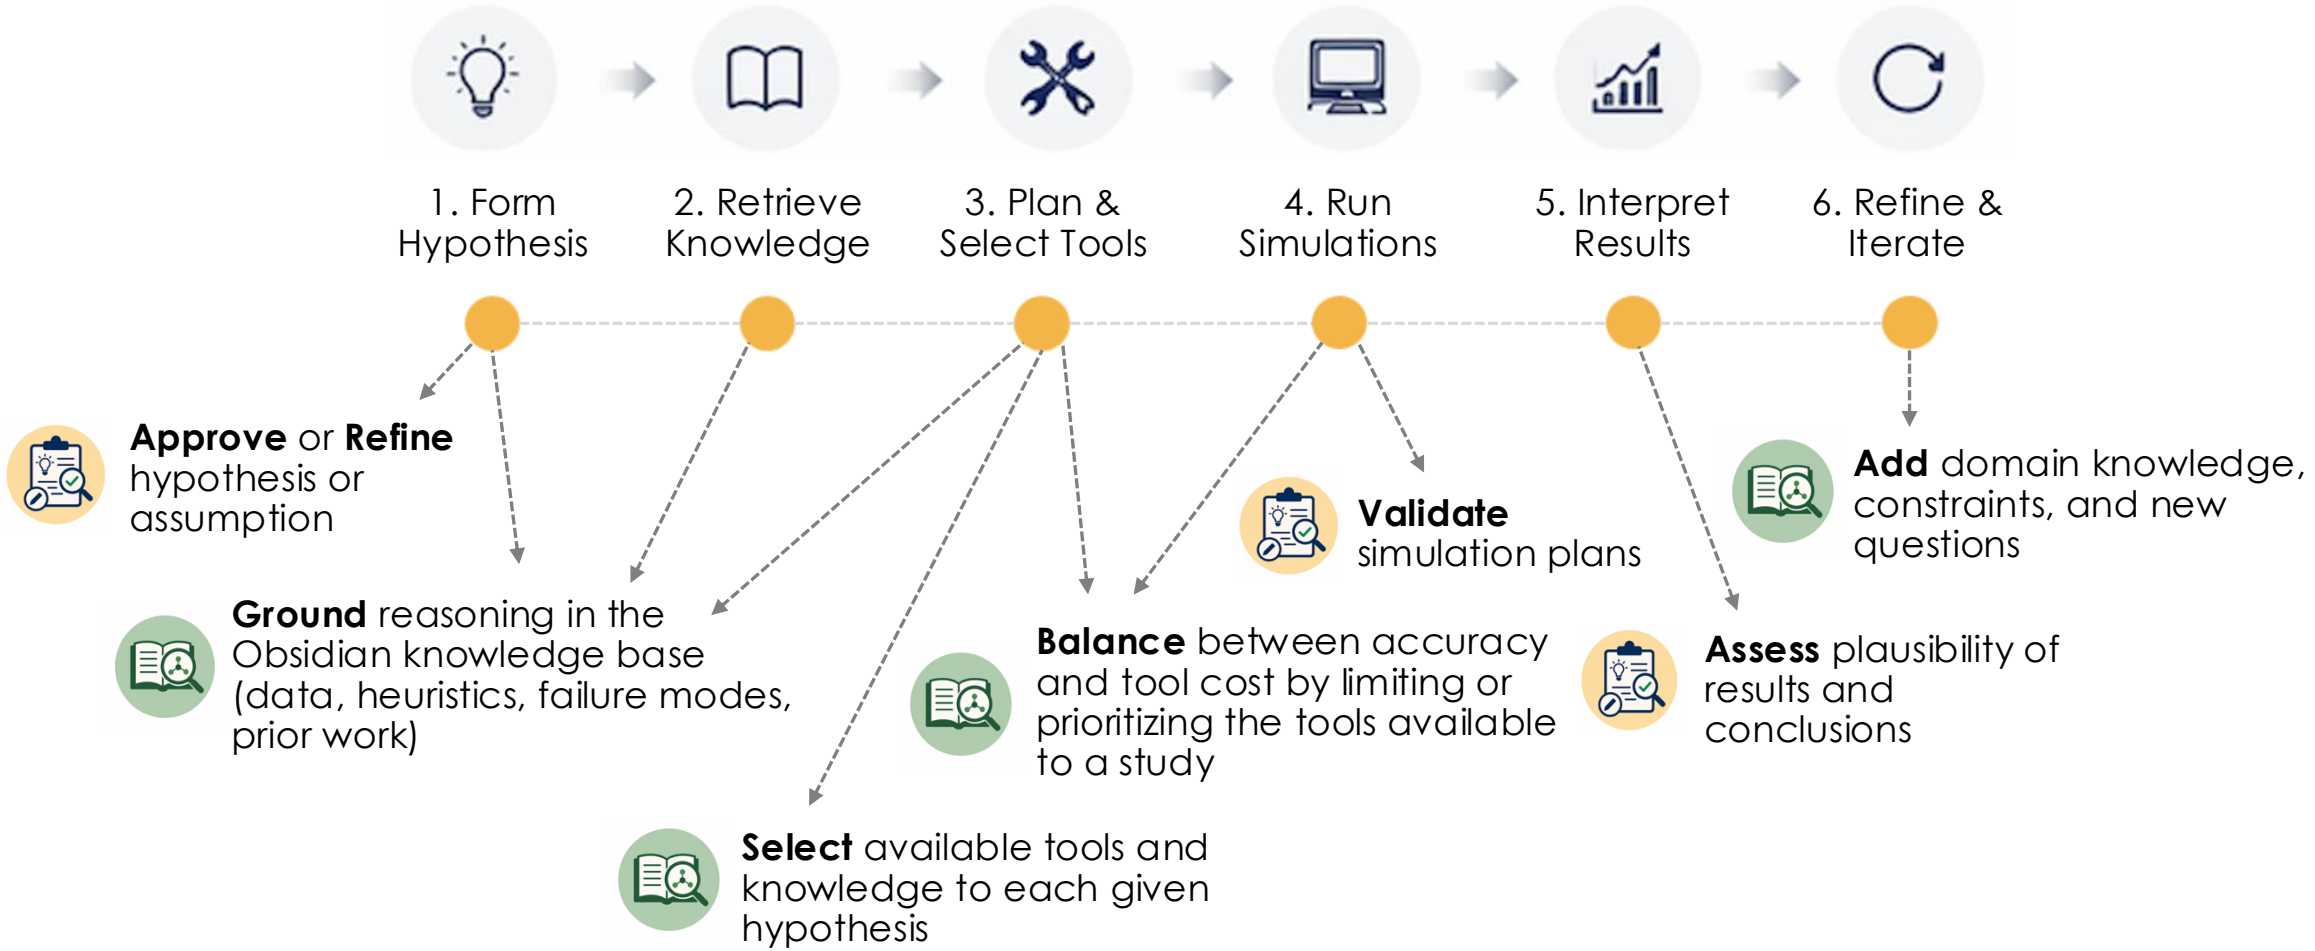

Human as Arbiter: Keeping Scientists in the Driver's Seat

Ensuring human judgment remains the ultimate authority on plausibility and knowledge

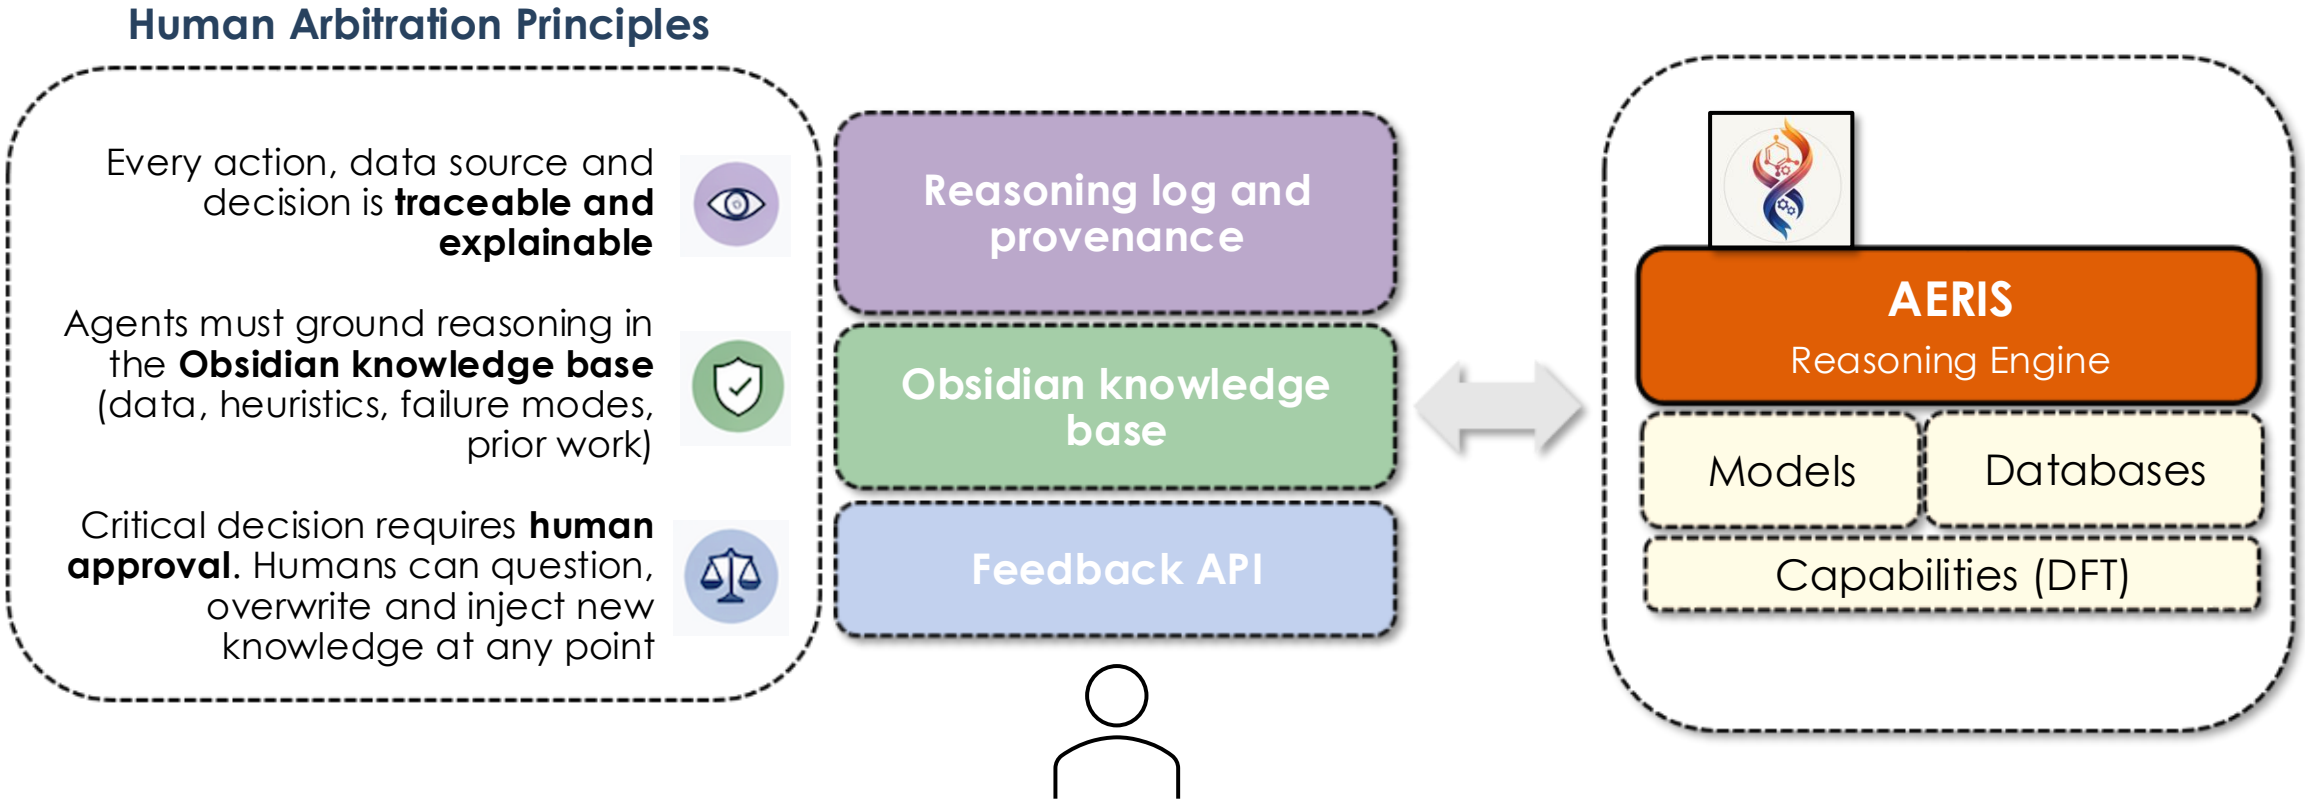

" The goal is not to slow the agent down-it is to ensure the science moves forward with transparency, rigor, and human judgment as the final authority "

Open Questions: When Agents Become Co-Investigators?

https://github.com/ORNL-Inria/AERIS
gainarua@ornl.gov
adornolopesd@ornl.gov

What counts as a "scientific decision"?

- Is selecting a model a decision?
- Is defining the search space a hypothesis?
- Is interpreting results reasoning or post-processing?

How do we evaluate agent-generated science?

- Accuracy vs scientific insight
- Agreement with known results vs novel discovery
- Benchmarks vs real scientific workflows

What does "co-PI" actually mean?

- Autonomy in decisions?
- Contribution to knowledge?
- Attribution / credit?
- At what point does assistance become collaboration?

Where should human oversight occur?

- At hypothesis generation?
- At tool selection?
- At result interpretation?
- How do we ensure agents amplify, not replace, human understanding.