Argonne NATIONAL LABORATORY

Agentic AI & Science: The Genesis Mission

Ian Foster

Argonne National Laboratory

The University of Chicago

Midwest RCD Annual Meeting, May 7, 2026

UChicago ▶ Argonne LLC

U.S. DEPARTMENT OF ENERGY

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

Remember the good old days?

- Another year, another 30% faster processors
- Another procurement cycle, another press release about “unprecedented computational capability”
- Another machine with 4× more CPUs and the same documentation copied forward from the last system
- Another parallel file system that would finally solve I/O bottlenecks forever
- Another heroic scaling plot showing 92% efficiency to 131,072 cores on a carefully chosen benchmark problem
- Another set of checkpoint/restart scripts held together with Bash, hope, and institutional memory
- Another reassuringly linear roadmap: more cores, more memory, more FLOPS, repeat
- And we should really do something about archival storage …

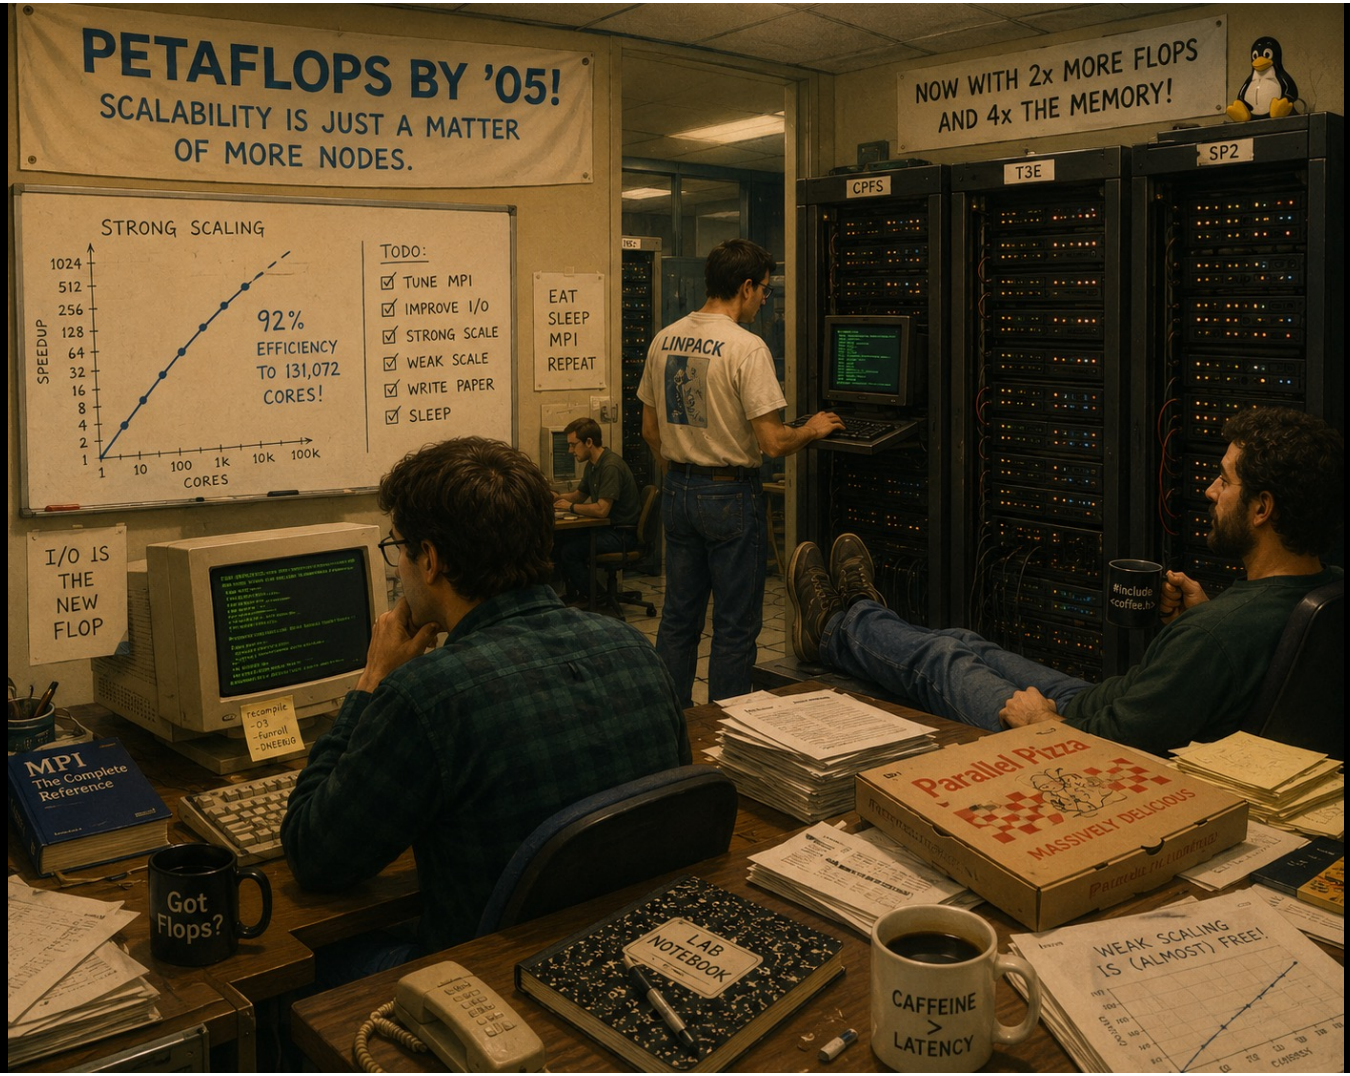
PETAFLOPS BY '05!
SCALABILITY IS JUST A MATTER OF MORE NODES.
NOW WITH 2x MORE FLOPS AND 4x THE MEMORY!
CPFS
T3E
SP2
STRONG SCALING
SPEEDUP
1024
512
256
128
64
32
16
8
4
2
1
1
10
100
1k
10k
100k
CORES
92% EFFICIENCY TO 131,072 CORES!
TODO:
TUNE MPI
IMPROVE I/O
STRONG SCALE
WEAK SCALE
WRITE PAPER
SLEEP
EAT SLEEP MPI REPEAT
LINPACK
I/O IS THE NEW FLOP
recompile -O3 -funroll -DNEEBUG
MPI The Complete Reference
Got Flops?
Parallel Pizza
MASSIVELY DELICIOUS
LAB NOTEBOOK
CAFFEINE > LATENCY
WEAK SCALING IS (ALMOST) FREE!

Now we live in interesting times

- HPC is no longer the domain of eccentric chemists hand-tuning Fortran kernels at 2AM while arguing about cache lines and basis sets.
- The users arrive with trillion-parameter models, venture capital, and the expectation that the machine should also write the paper.
- The old measure of success was a simulation finishing before the allocation expired; now it is whether the model develops “emergent reasoning.”
- We used to optimize floating-point throughput; now we optimize GPU supply chains, power contracts, and human patience.
- Scientific computing once meant understanding the physics well enough to model it. Now it means understanding the model well enough to trust its physics.
- The old systems solved equations. The new systems generate hypotheses, code, experimental plans, and occasionally entirely fictional citations.
- HPC has shifted from accelerating science to participating in it — and nobody is entirely sure what that means yet.

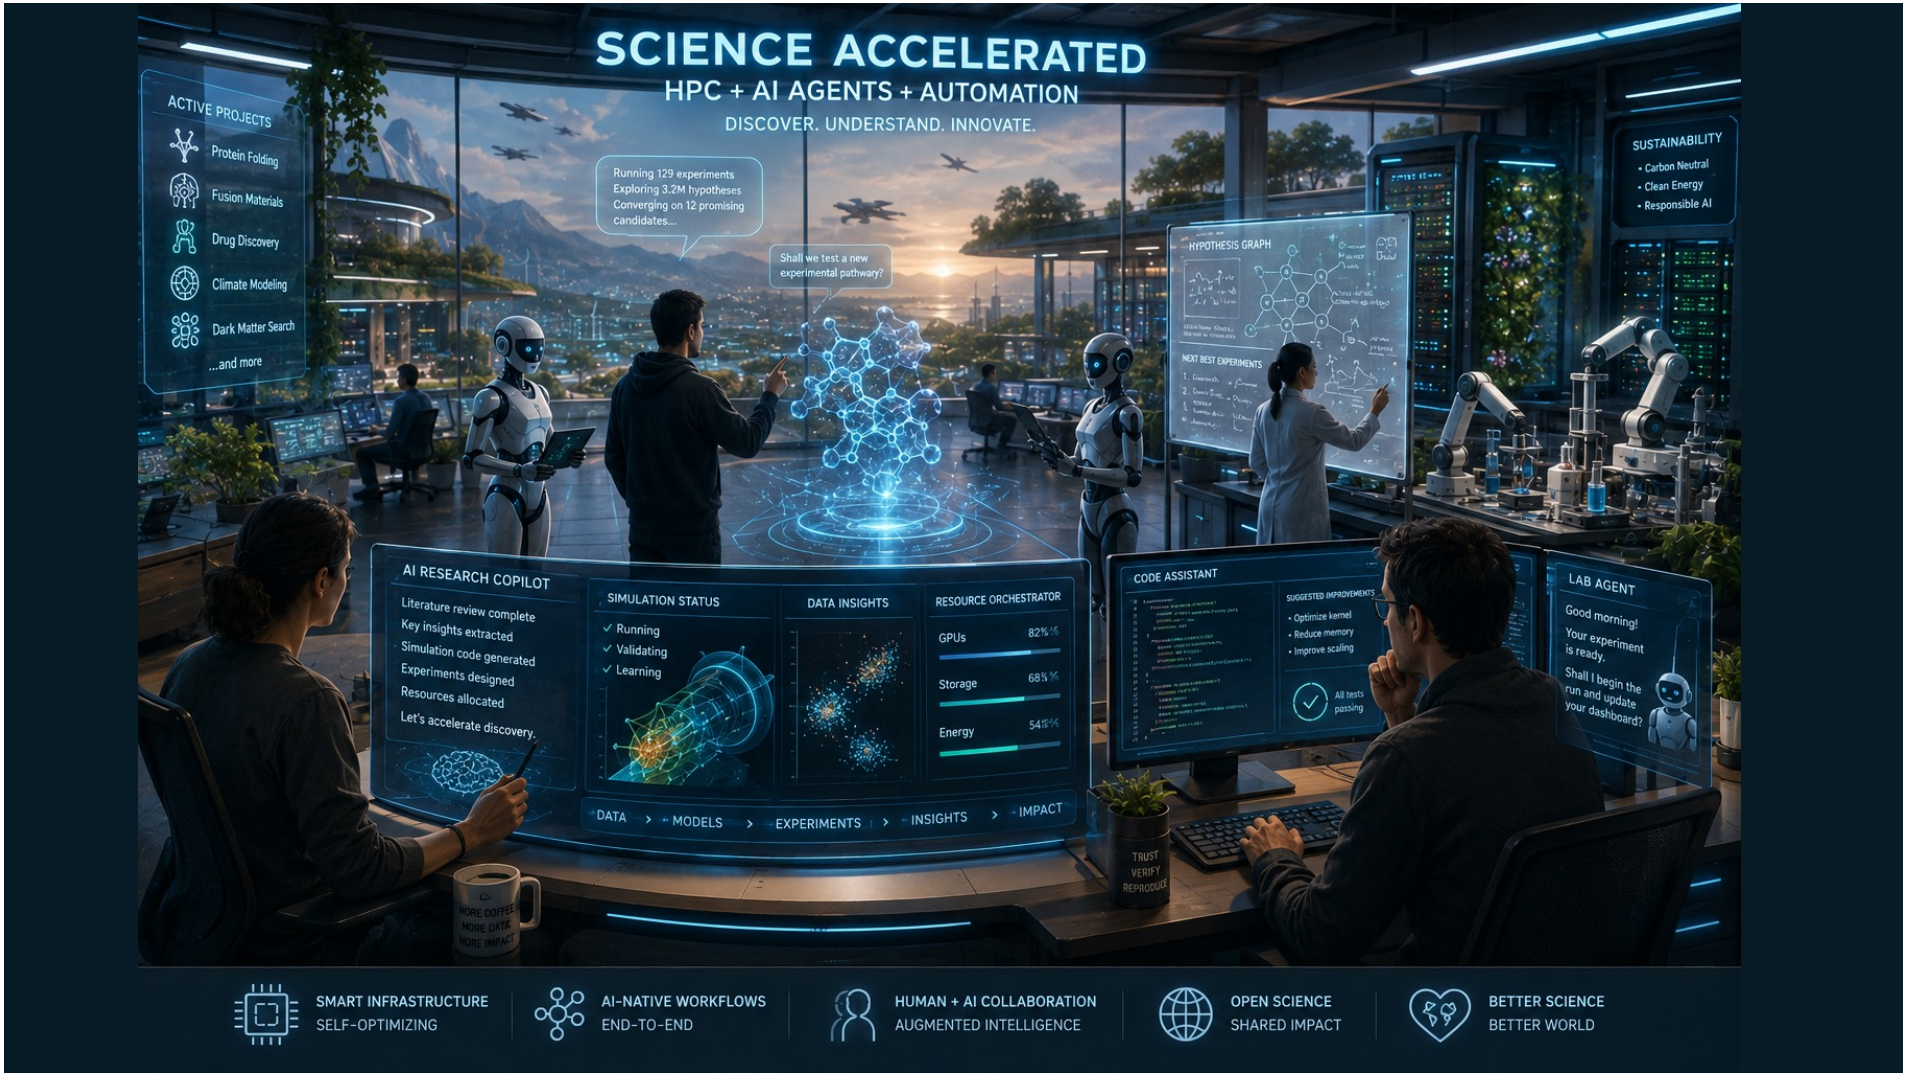
SCIENCE ACCELERATED
HPC + AI AGENTS + AUTOMATION
DISCOVER. UNDERSTAND. INNOVATE.
ACTIVE PROJECTS
Protein Folding
Fusion Materials
Drug Discovery
Climate Modeling
Dark Matter Search
...and more
Running 129 experiments
Exploring 3.2M hypotheses
Converging on 12 promising candidates...
Shall we test a new experimental pathway?
HYPOTHESIS GRAPH
NEXT BEST EXPERIMENTS
SUSTAINABILITY
• Carbon Neutral
• Clean Energy
• Responsible AI
AI RESEARCH COPILOT
Literature review complete
Key insights extracted
Simulation code generated
Experiments designed
Resources allocated
Let's accelerate discovery.
SIMULATION STATUS
✓ Running
✓ Validating
✓ Learning
DATA INSIGHTS
RESOURCE ORCHESTRATOR
GPUs 82%
Storage 68%
Energy 54%
CODE ASSISTANT
SUGGESTED IMPROVEMENTS
• Optimize kernel
• Reduce memory
• Improve scaling
All tests passing
LAB AGENT
Good morning!
Your experiment is ready.
Shall I begin the run and update your dashboard?
DATA > MODELS > EXPERIMENTS > INSIGHTS > IMPACT
MORE COFFEE
MORE DATA
MORE IMPACT
TRUST
VERIFY
REPRODUCE
SMART INFRASTRUCTURE
SELF-OPTIMIZING
AI-NATIVE WORKFLOWS
END-TO-END
HUMAN + AI COLLABORATION
AUGMENTED INTELLIGENCE
OPEN SCIENCE
SHARED IMPACT
BETTER SCIENCE
BETTER WORLD

Graduate-Level Google-Proof Q&A test (GPQA), Diamond problems

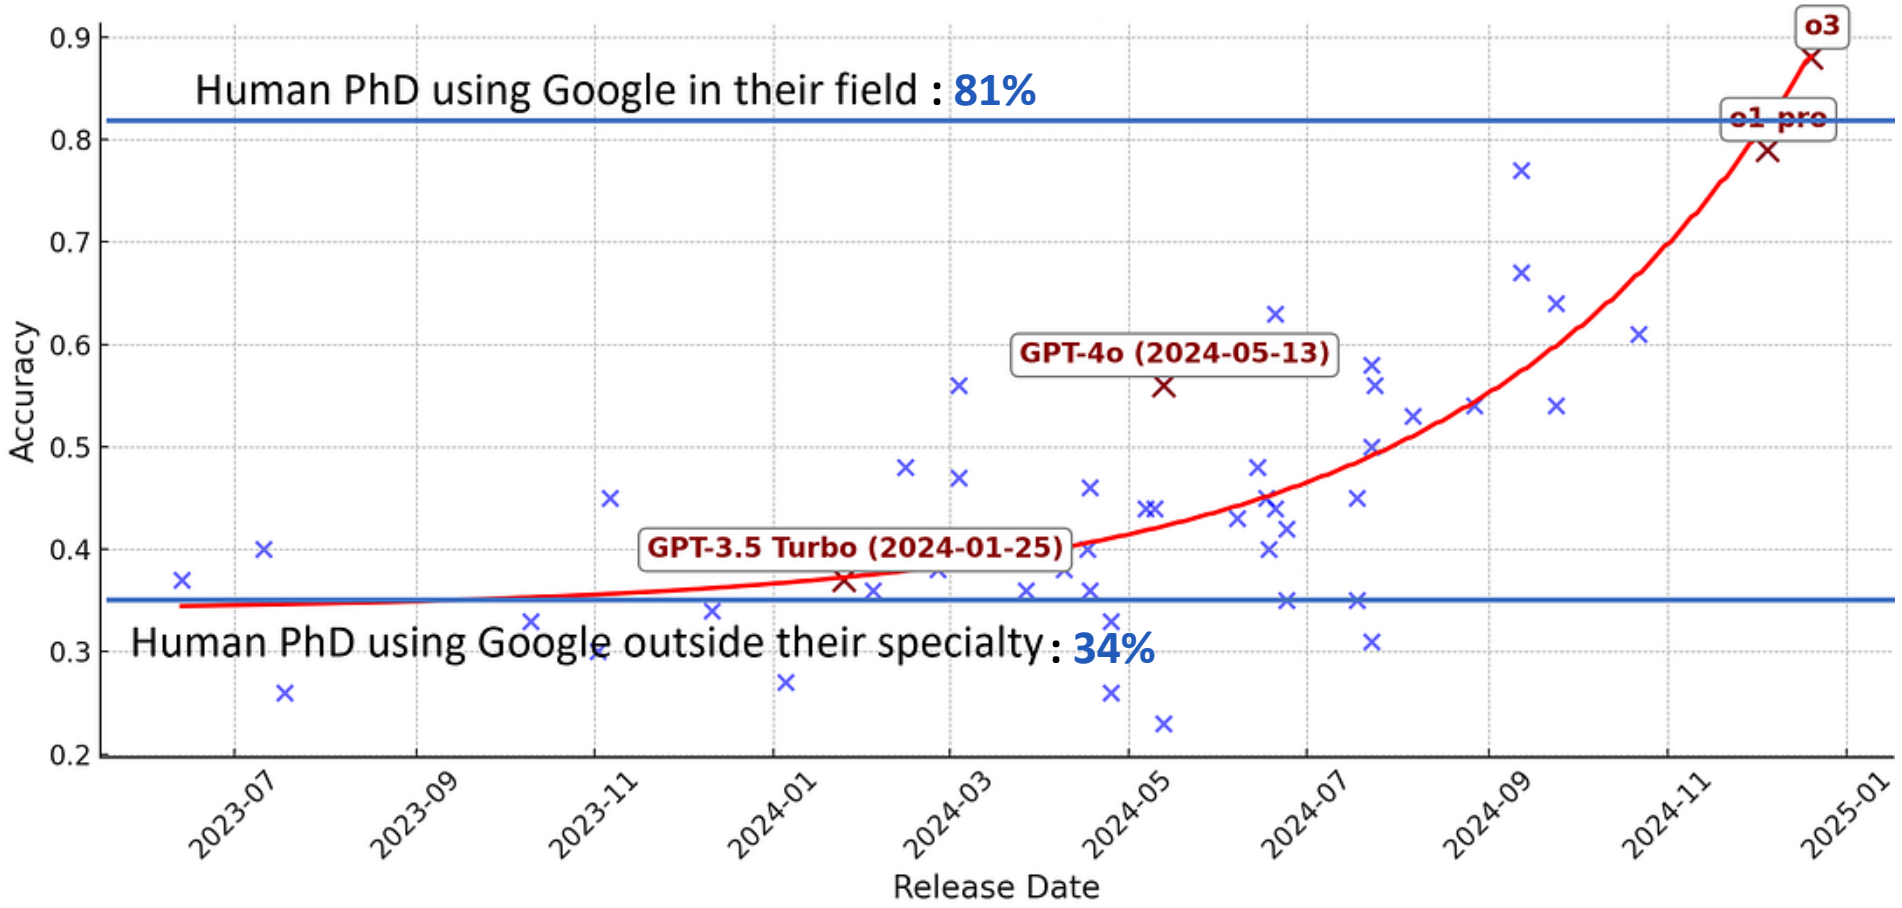

https://arxiv.org/pdf/2311.12022

https://epoch.ai/data/ai-benchmarking-dashboard

Traditional ML

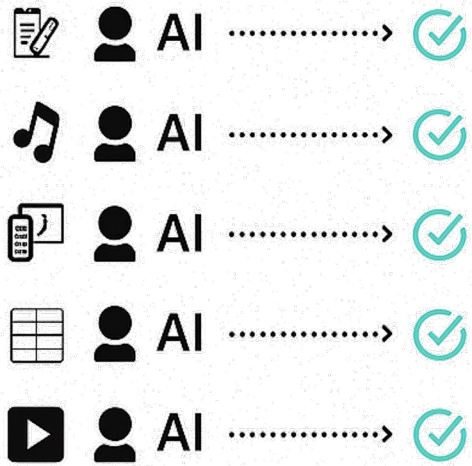

- Individual siloed models
- Require task-specific training
- Lots of human supervised training

Foundation models
(e.g., Large Language Models: LLMs)

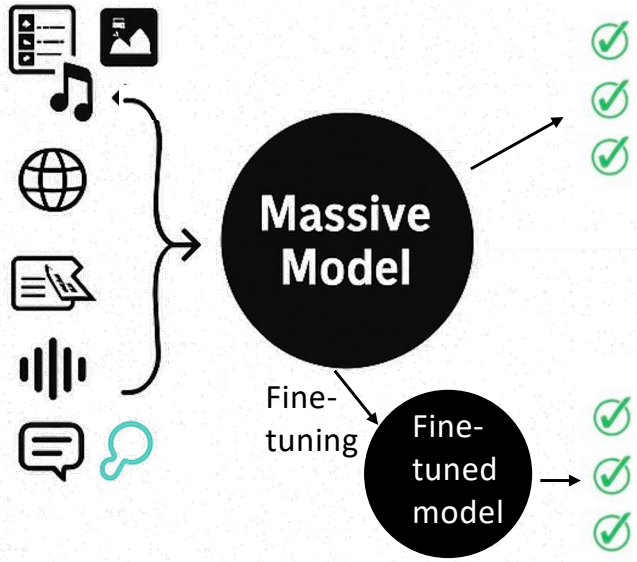

- Massive multi-modal model
- Adapted with minimal training
- Pre-trained unsupervised learning

Reasoning models (RMs)

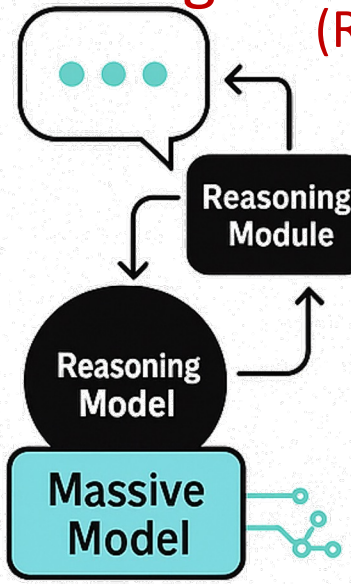

- Deliberative multi-step logic
- Self-consistency checks
- Slower but more accurate inference

FMs/RMs are general-purpose technologies

Humans engage FMs for many purposes, e.g.:

- Analyze **knowledge**
- Define & evaluate **hypotheses**
- **Define protocols** to test
- **Select data** to use or request
- **Choose tools** (e.g., simulators, instruments, computers)
- **Define actions** (e.g., launch job, run query, trigger experiment)
- **Evaluate outputs**
- **Propose next steps**

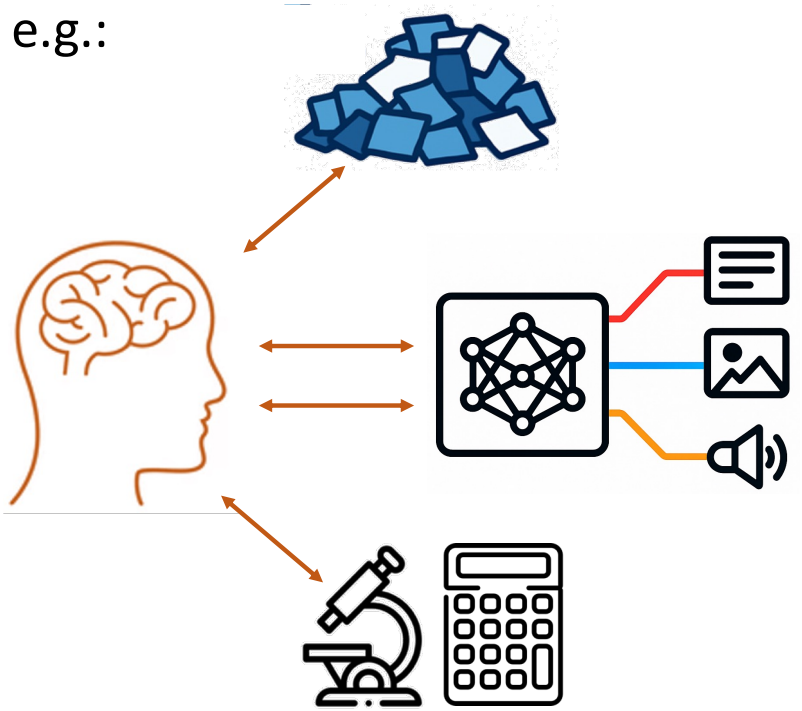

Human decision-making increasingly becomes the bottleneck

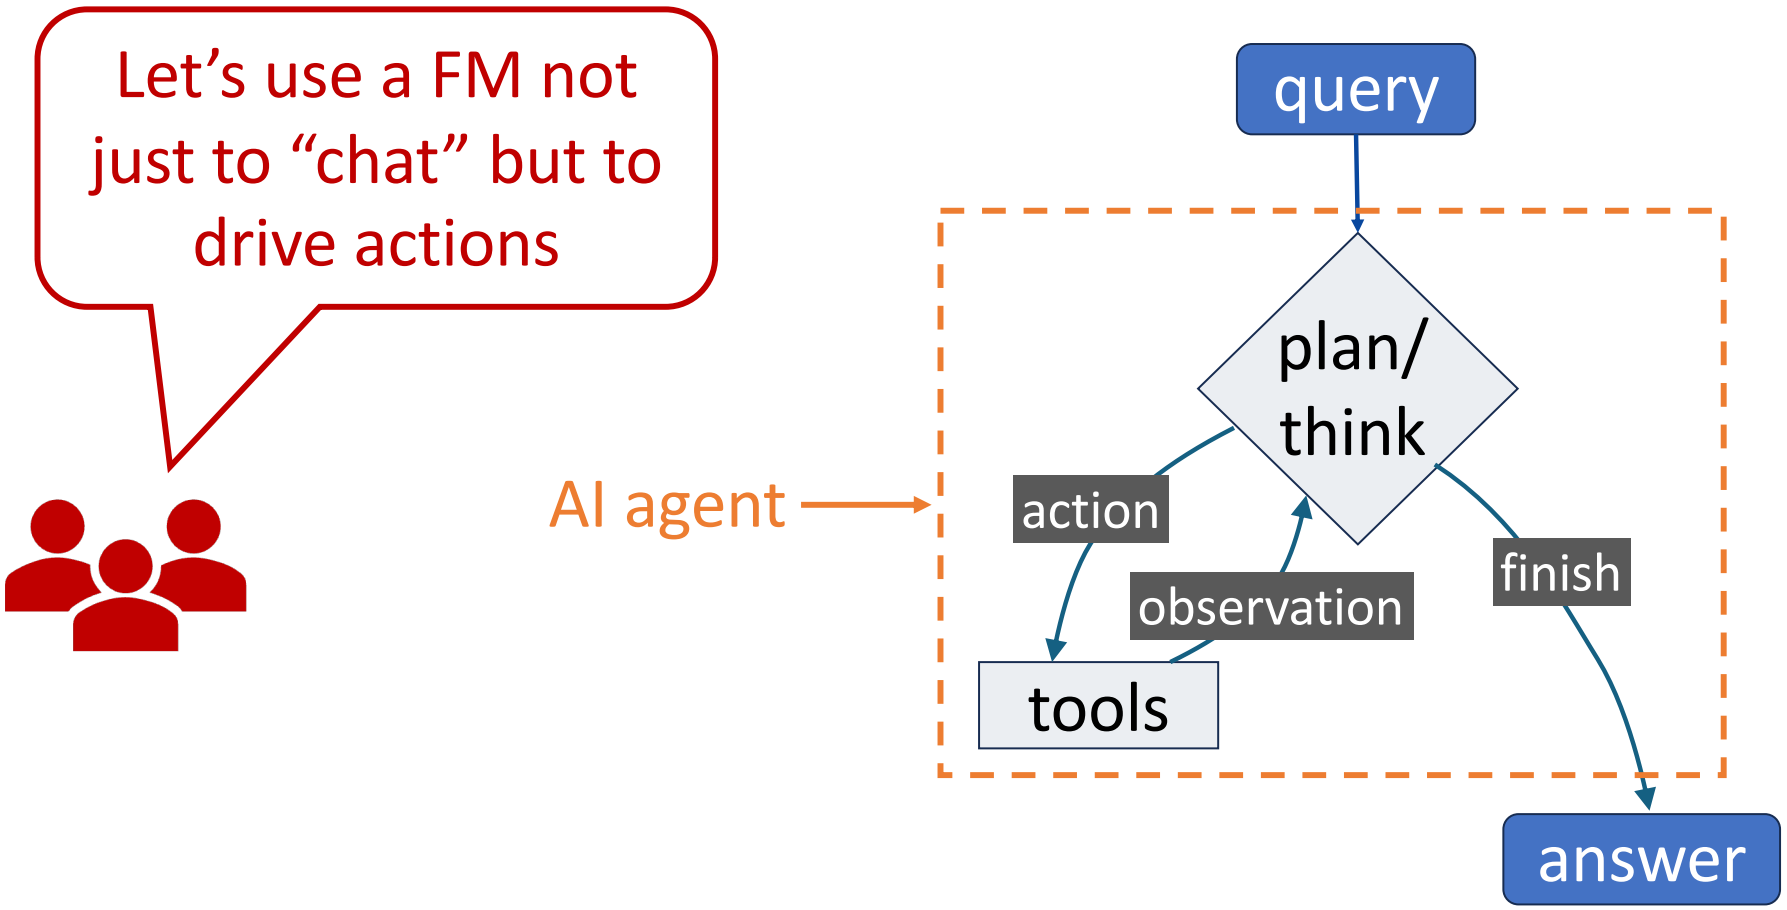
Let’s use a FM not just to “chat” but to drive actions
AI agent
query
plan/ think
action
observation
finish
tools
answer

An agent is a persistent, stateful process that acts on behalf of a user or system

An agent may:

- **Observe** inputs or events [or user query]
- **Plan** (decide on) actions using a policy (rules or FM)
- **Act**: Execute tools or call other agents
- **Learn**: Update state to adapt over time

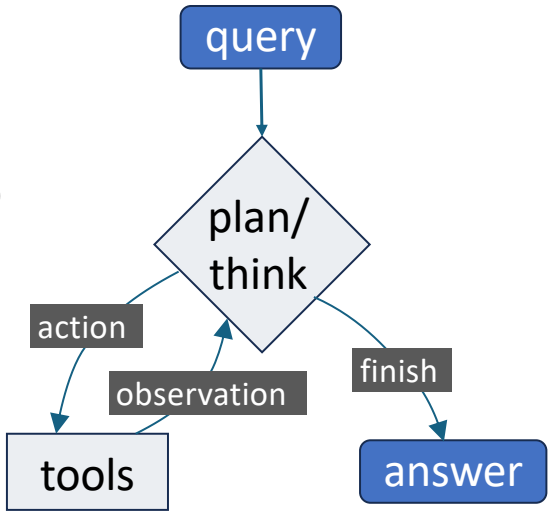

We can think of an agent as a [scientific] assistant that can reason, act, and coordinate on our behalf

An agentic architecture for the design of antimicrobial peptides

An antimicrobial peptide (AMP) is a short (typically 12 to 50 amino acid) molecule that can target and kill viruses, bacteria, fungi, and other pathogens

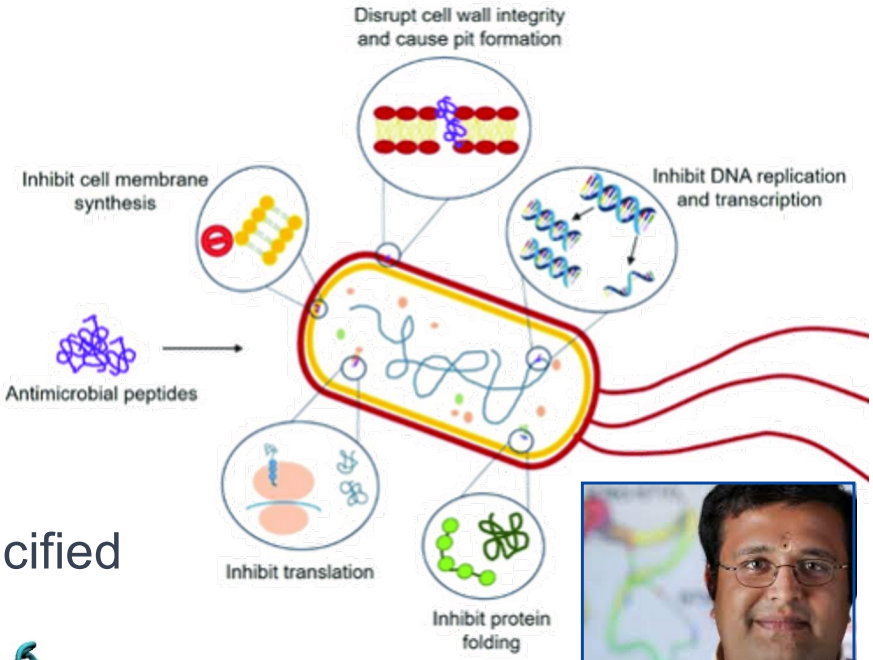

Challenge: Design an AMP that can kill specified bacterial strains without harming host cells

With 20 possible amino acids, there are $20^{20} = 10^{26}$ AMPs of length 20

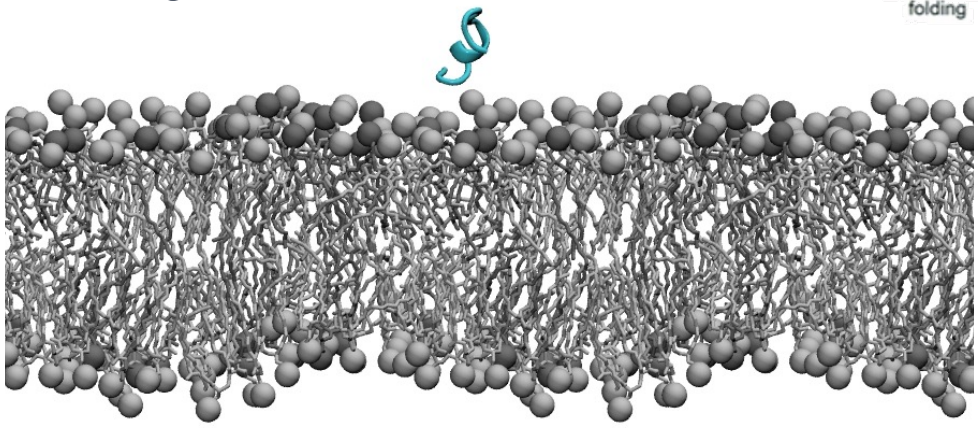

Arvind Ramanathan

An agentic architecture for the design of antimicrobial peptides

An antimicrobial peptide (AMP) is a short (typically 12 to 50 amino acid) molecule that can target and kill viruses, bacteria, fungi, and other pathogens

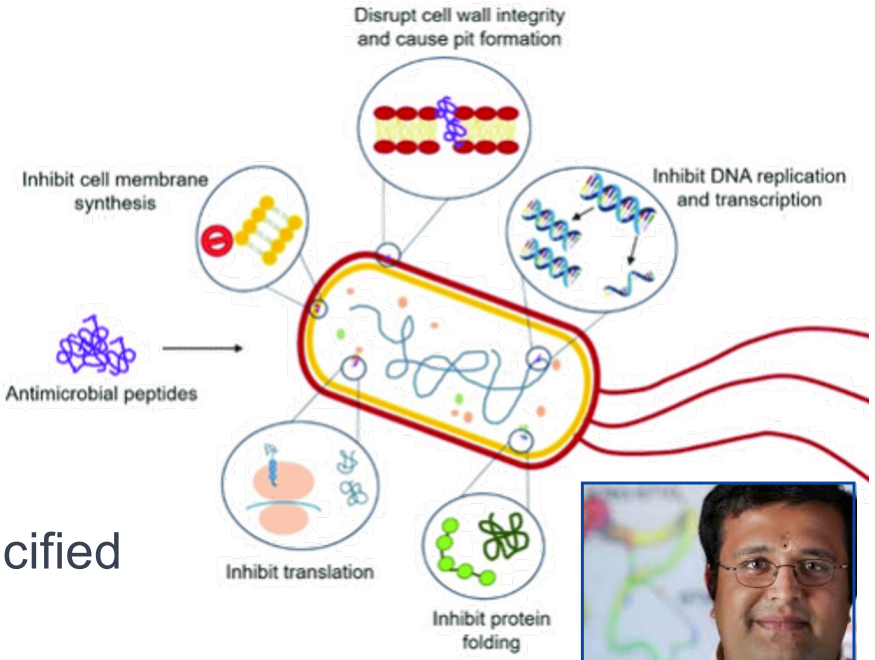

Challenge: Design an AMP that can kill specified bacterial strains without harming host cells

Arvind Ramanathan

A rational design approach might combine knowledge of bacterial cell membrane composition and structure, AMP molecular and structural properties, host cell membrane characteristics and intracellular pathways—knowledge that may be gained by **database/literature search**, **simulation**, **experiment**

Example: A **peptide expert**

(Prototyped with PubMed and ChatGPT)

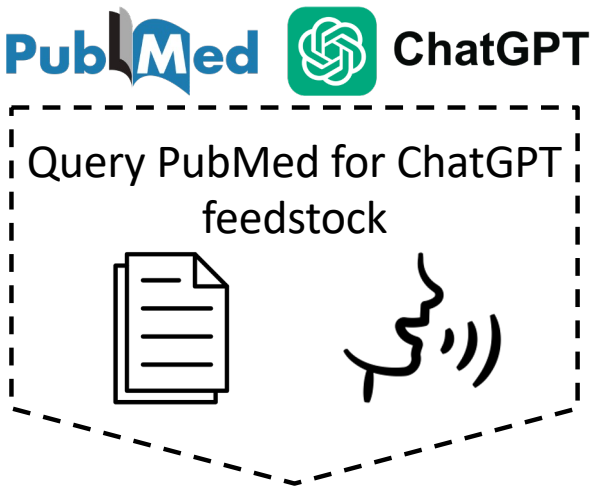

Retrieve abstracts **A** from PubMed that reference specified **peptide**

Use ChatGPT to build hypotheses via retrieval-augmented generation: e.g.:

"Given **A**, on which organism is {**peptide**} acting?"

Evaluate hypotheses and update query and hypothesis generation policies

Arvind Ramanathan, Priyanka Setty, et al.

Outline of a peptide agent implementation

```python
# Persistent agent state
Hypotheses = {}         # hypothesis_id -> hypothesis object
EvidenceLedger = {}     # hypothesis_id -> list of evidence events
RetrievalPolicy = {}    # learned biases over queries/sources
AgentConfig = {
    "confidence_threshold": 0.8,
    "retirement_threshold": -0.5
}

while True:

    # 1. Select target peptide
    peptide = select_peptide(Hypotheses)

    # 2. Retrieve new evidence
    query = build_pubmed_query(peptide, RetrievalPolicy)
    abstracts = retrieve_abstracts(query)

    # 3. Generate new hypotheses (if needed)
    if not has_active_hypothesis(peptide, Hypotheses):
        new_hypotheses = generate_hypotheses(
            abstracts,
            prompt = (
                "Given abstracts A, on which organism is "
                f"{peptide} acting?"
            )
        )
        for H in new_hypotheses:
            Hypotheses[H.id] = H
            EvidenceLedger[H.id] = []

    # 4. Evaluate hypotheses against new evidence
    for H in active_hypotheses(peptide, Hypotheses):

        assessment = evaluate_hypothesis(H, abstracts)
        # assessment ∈ {supports, contradicts, inconclusive}
        # with confidence score

        EvidenceLedger[H.id].append({
            "abstracts": abstracts,
            "assessment": assessment.type,
            "strength": assessment.strength,
            "timestamp": now()
        })

        # 5. Update hypothesis confidence
        H.confidence = update_confidence(
            H.confidence,
            assessment
        )
```

Tool call

“Think”

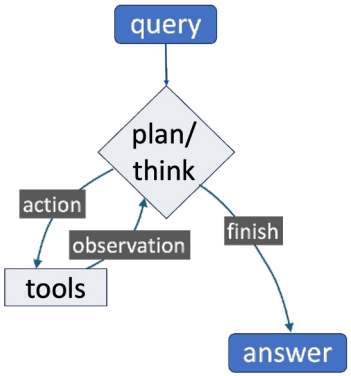

Learn

Finish

```python
# 6. Prune or retire hypotheses
for H in Hypotheses.values():
    if H.confidence < AgentConfig["retirement_threshold"]:
        retire_hypothesis(H)

# 7. Adapt retrieval policy
RetrievalPolicy = update_retrieval_policy(
    RetrievalPolicy,
    EvidenceLedger
)

# 8. Self-reflection and gap analysis (periodic)
if time_for_reflection():
    gaps = identify_knowledge_gaps(Hypotheses)
    RetrievalPolicy = bias_toward_gaps(
        RetrievalPolicy,
        gaps
    )

# 9. Optional human-in-the-loop
if uncertainty_high(Hypotheses):
    request_human_feedback(Hypotheses)

# 10. Termination or sleep
if stopping_condition_met(Hypotheses):
    break
else:
    sleep()
```

Define other agents, also FM-powered

PubMed ChatGPT

Query PubMed for ChatGPT feedstock

Protein BLAST protein ▶ protein AlphaFold

Align proteins, predict structure, rank results

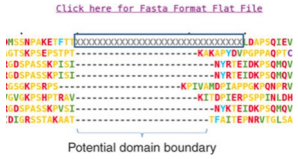

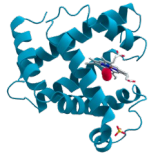

BV-BRC NCBI

Evaluate structures and filter results

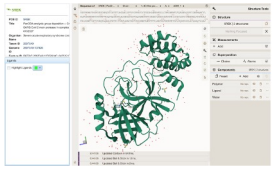

Arvind Ramanathan, Priyanka Setty, et al.

Link agents to construct an application

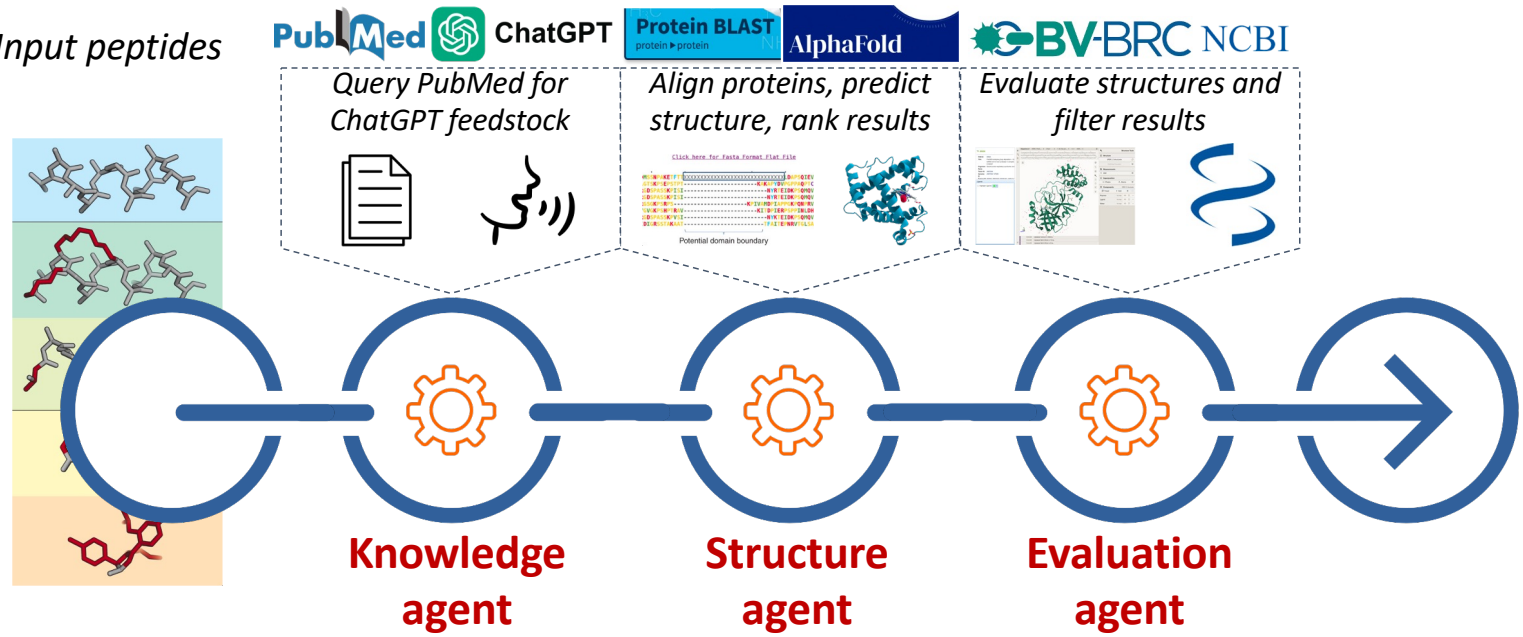

We use **Academy** to create and manage individual agents, which query databases, retrieve data, run simulations, run experiments, etc.

Arvind Ramanathan, Priyanka Setty, et al.

https://academy-agents.org

Link with HPC for **computational evaluation**

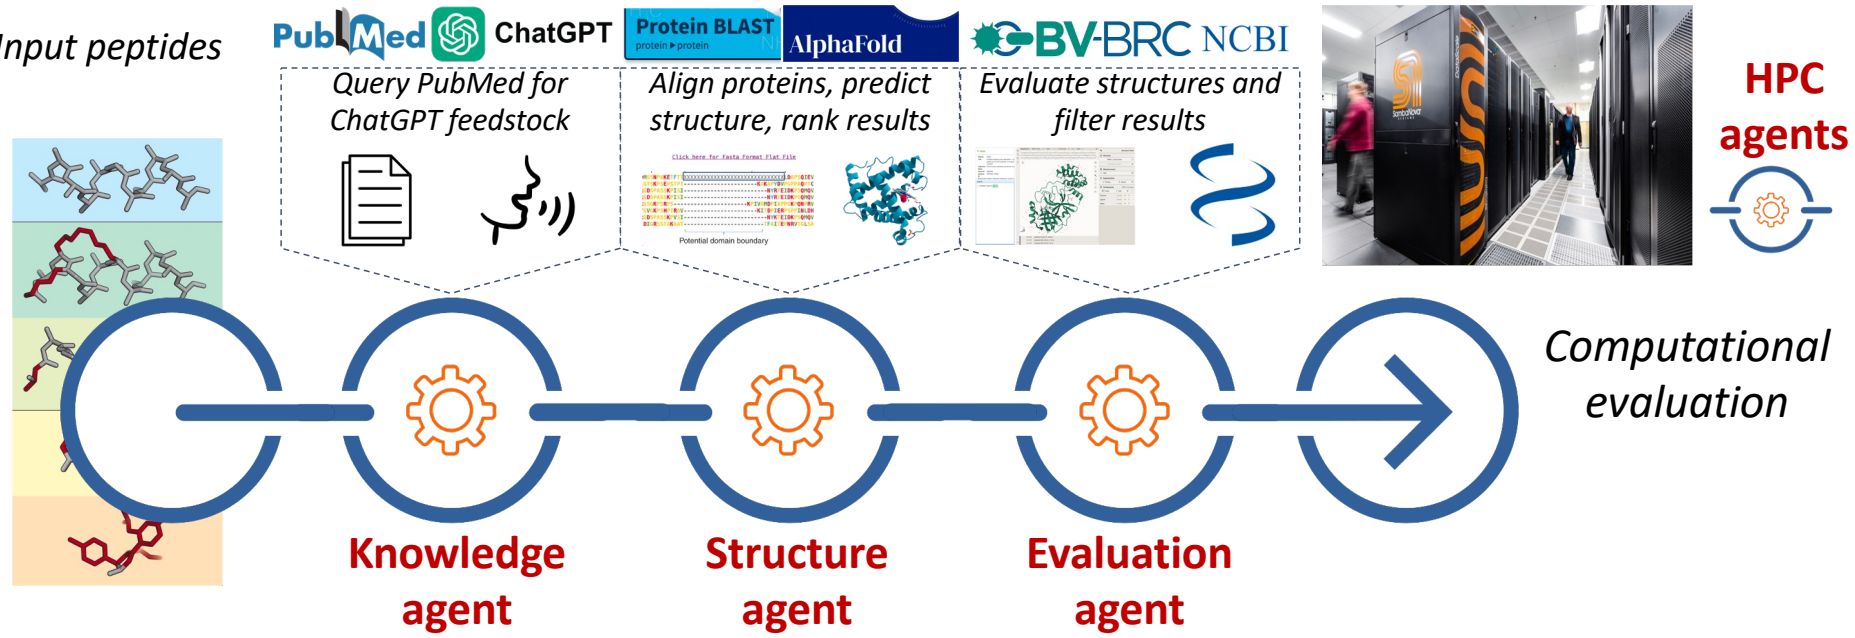

Arvind Ramanathan, Priyanka Setty, et al.

https://academy-agents.org

Link with self-driving labs for **experimental evaluation**

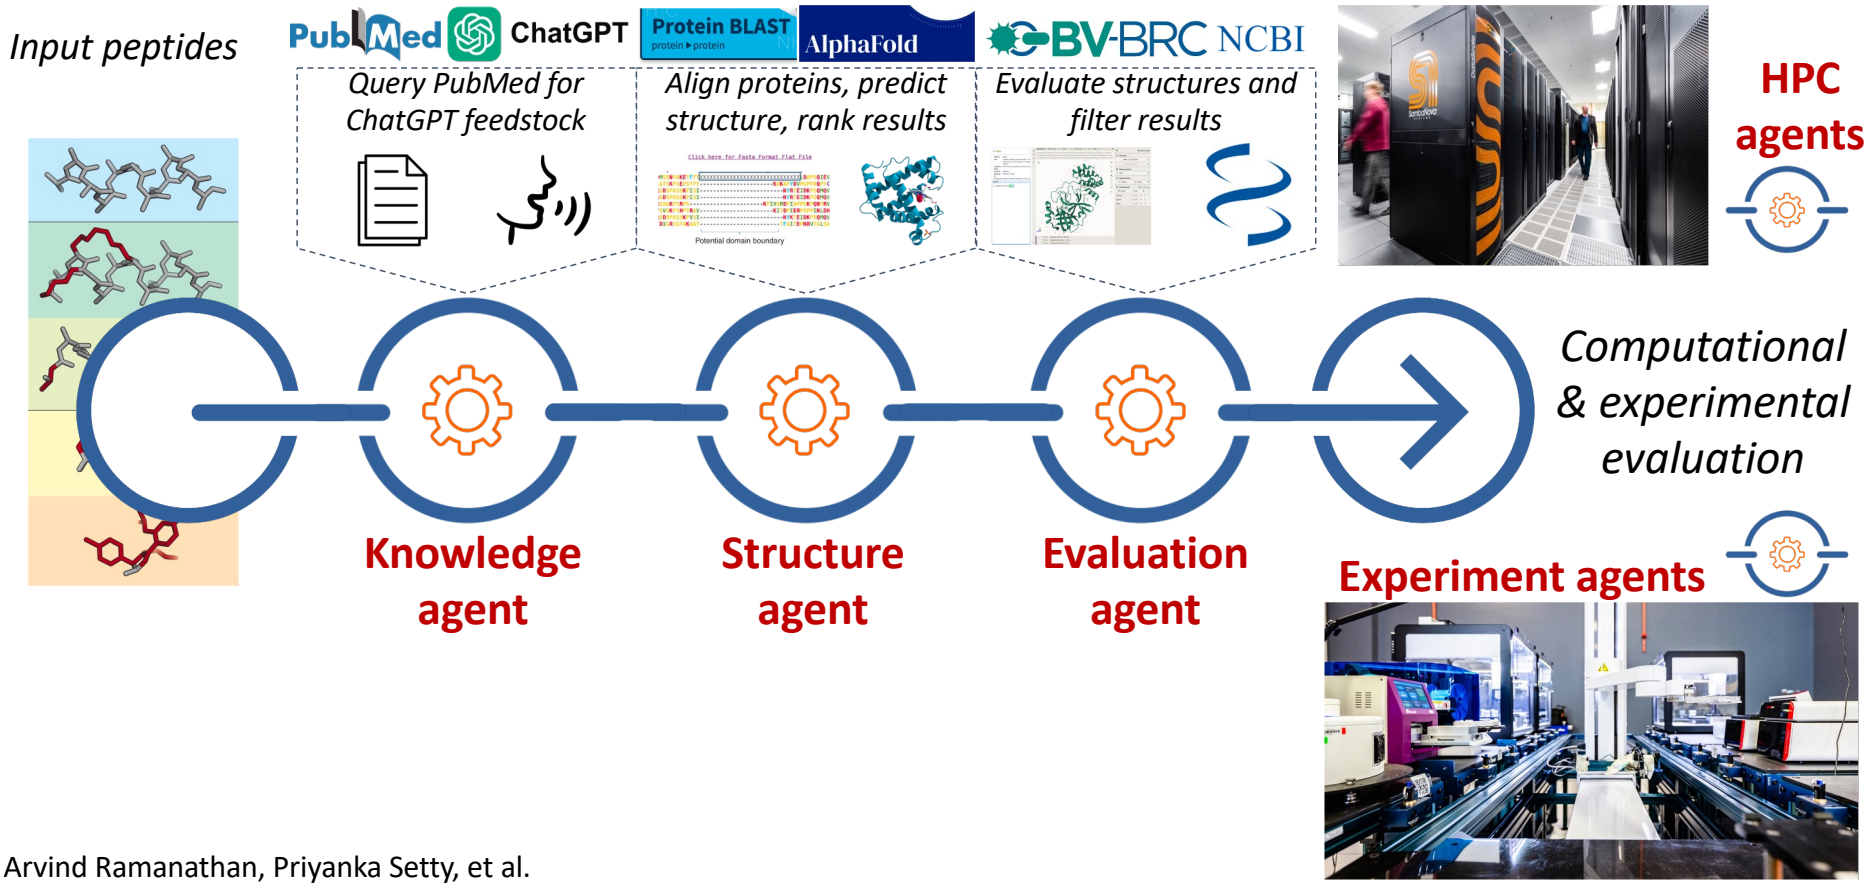

Arvind Ramanathan, Priyanka Setty, et al.

Close the loop for autonomous discovery

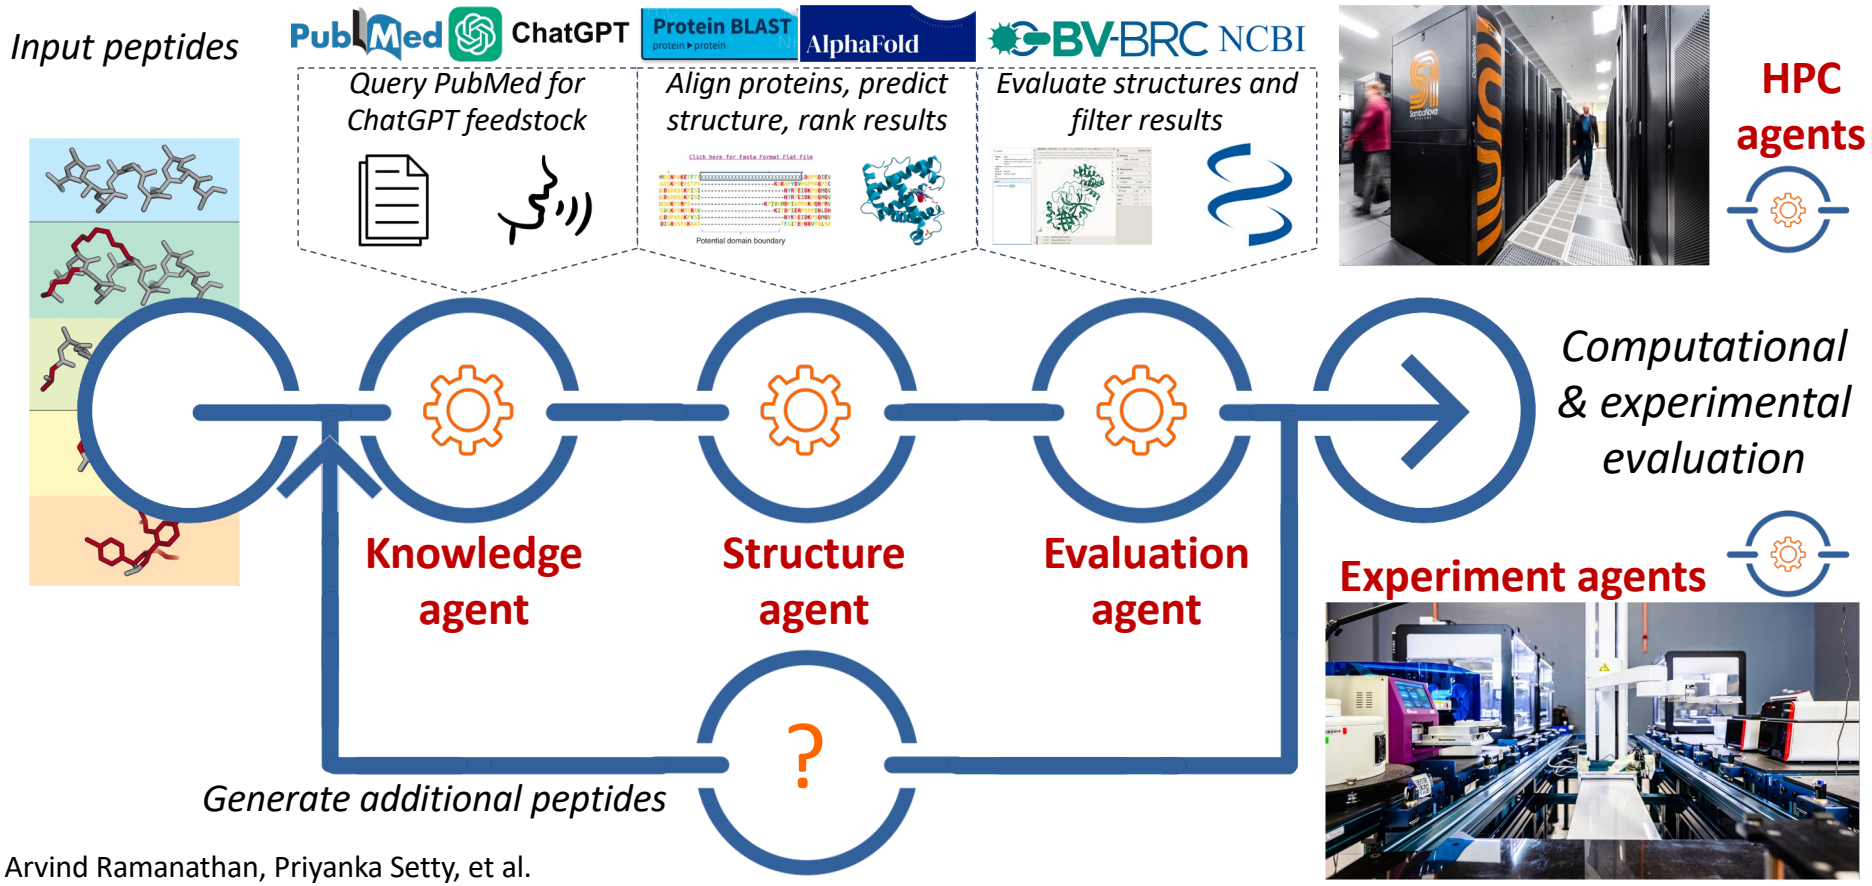

Arvind Ramanathan, Priyanka Setty, et al.

Observations

1) Everyone wants agents. How will we support and manage them (the users and the agents)?

2) Scientific agents need to interact with cyberinfrastructure, securely, reliably, and incessantly. How?

3) Agents will replace people as the dominant users of infrastructure. What does that mean for resource demands and policies?

Agentic middleware: Scope and challenges

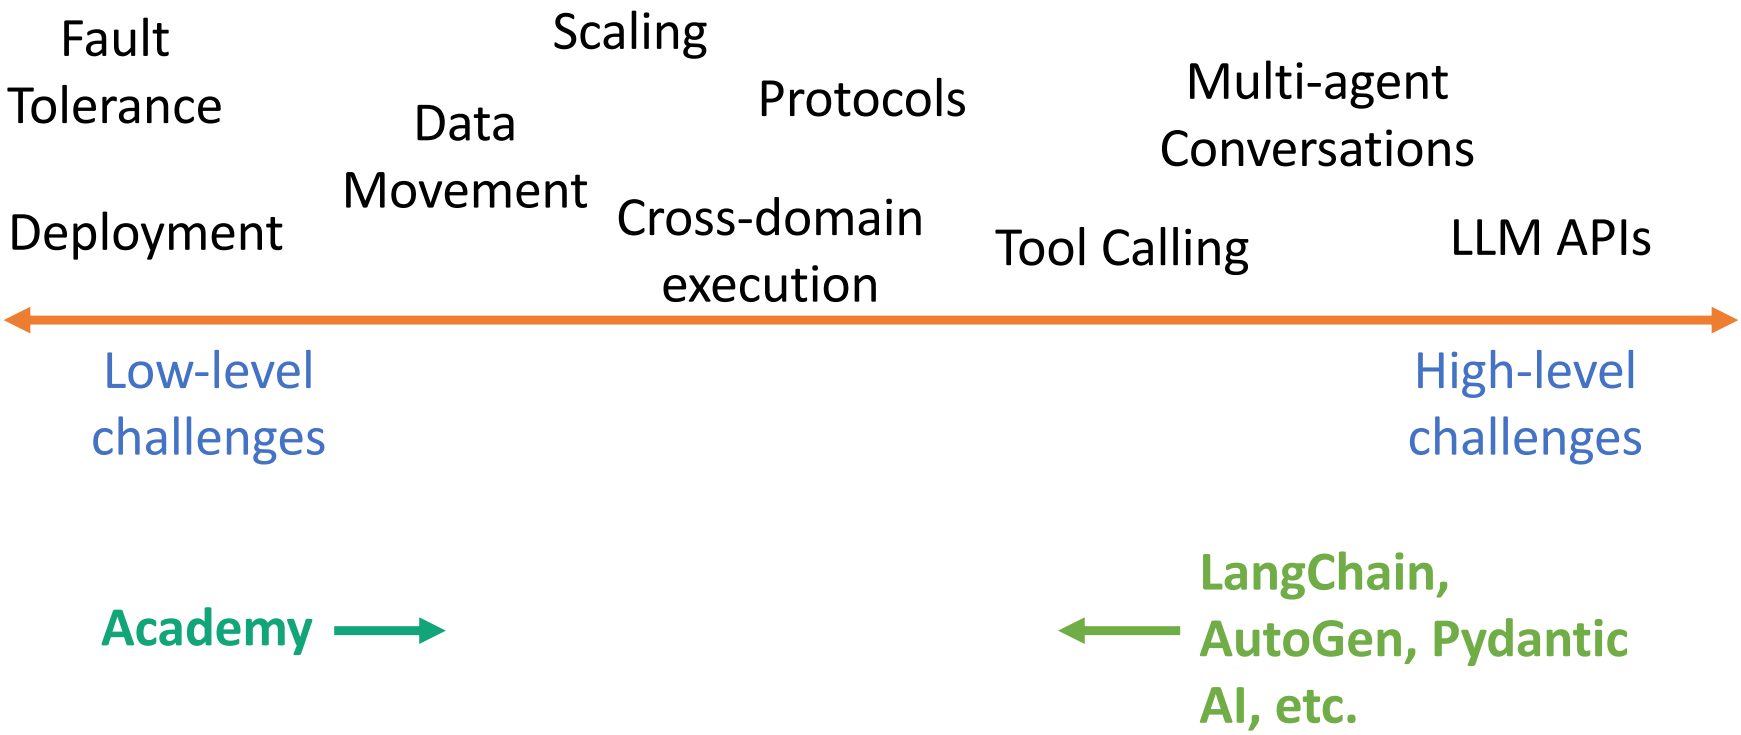

Exploring agentic middleware: **Academy**

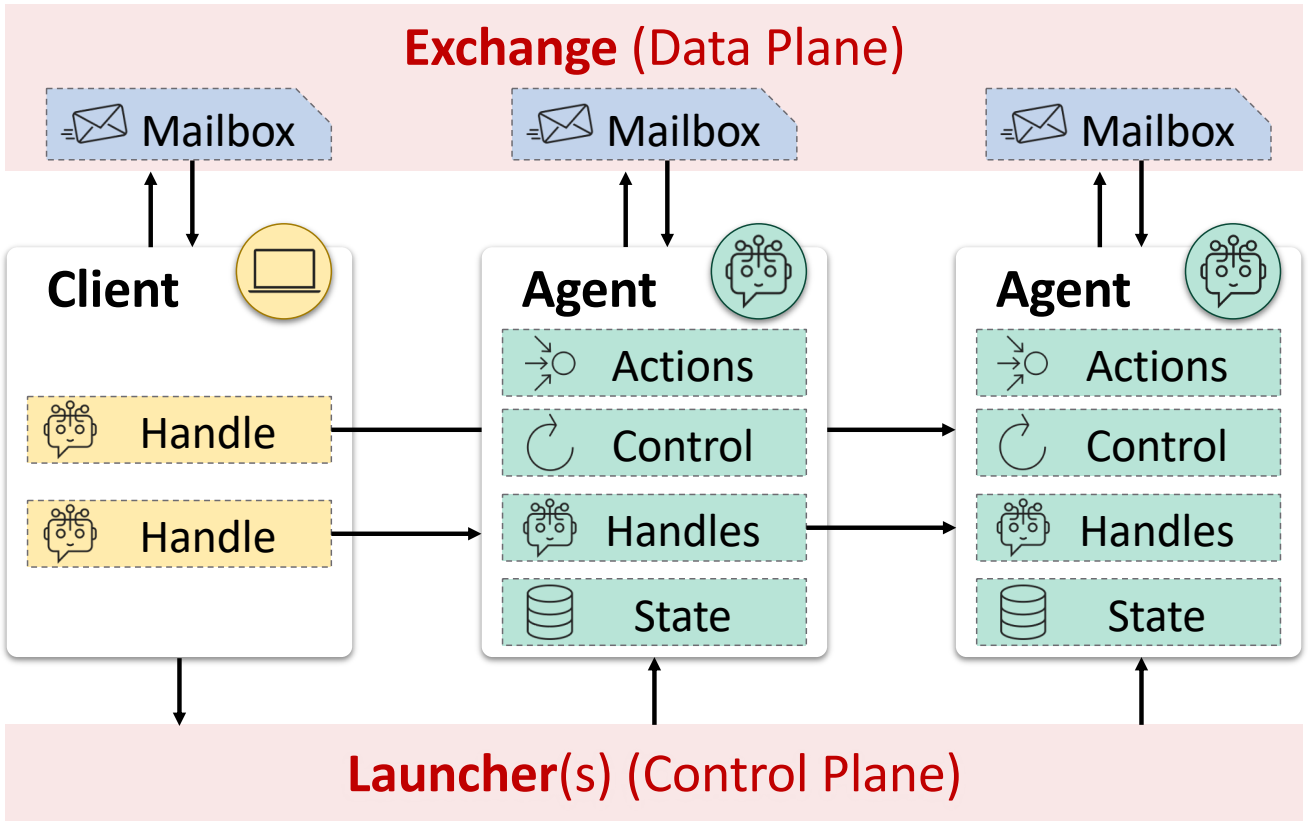

Dr. Greg Pauloski

Dr. Kyle Chard

Alok Kamatar

https://academy-agents.org

Agents defined by a **behavior**

Clients & other agents can request **actions**

```python
import asyncio
from academy.behavior import Behavior, action, loop

class Example(Behavior):
    def __init__(self) -> None:
        self.count = 0  # State stored as attributes

    @action
    async def square(self, value: float) -> float:
        return value**2

    @loop
    async def count(self, shutdown: asyncio.Event):
        while not shutdown.is_set():
            self.count += 1
            asyncio.sleep(1)
```

Instance of a behavior is **state**

Control loops for autonomous behavior

https://docs.academy-agents.org/latest/get-started/

Agentified MOF design application: Execution trace

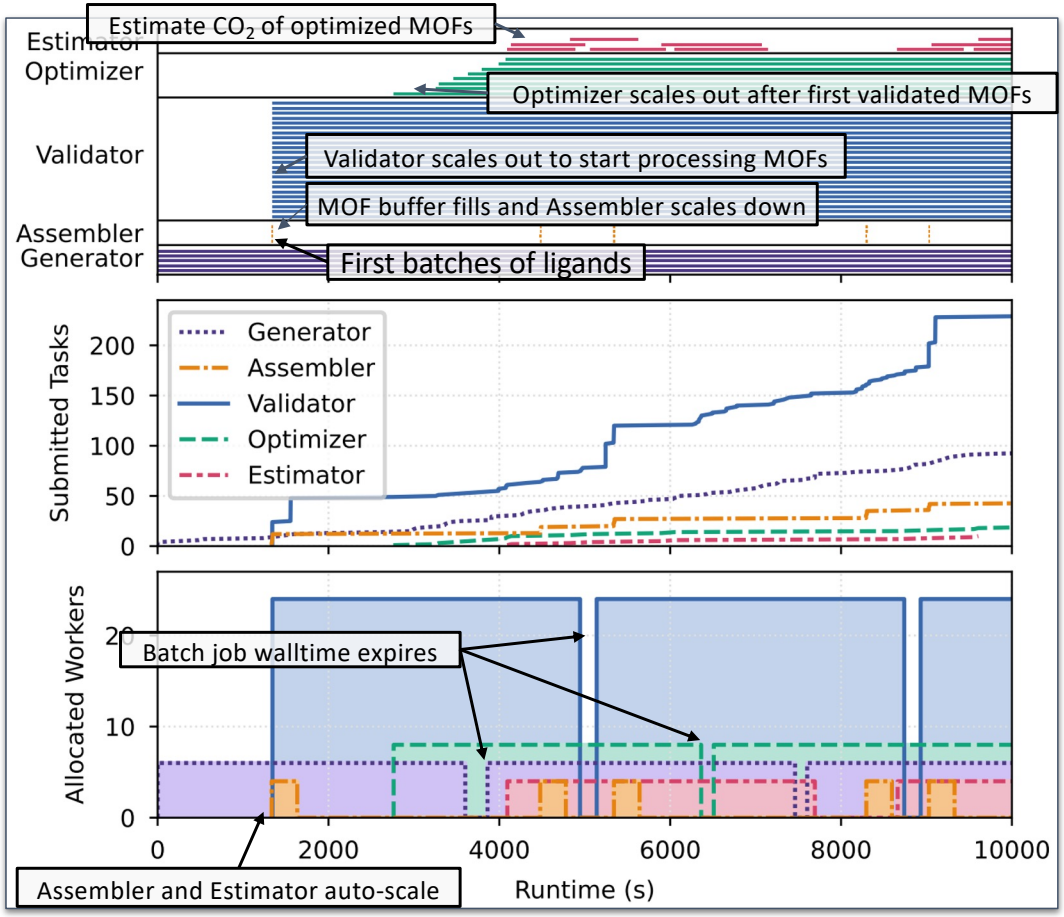

Benefits of agentic model:

- **Placement:** Move agents to resources
- **Separation of concerns:** Resource acquisition & scaling based on local workload
- **Loose coupling:** Swap agents, integrate new agents (e.g., selff-driving labs)
- **Shared agents:** Multiple workflows can share agents (microservice-like)

The Genesis Mission

A National Strategy to Accelerate Science Through Artificial Intelligence.

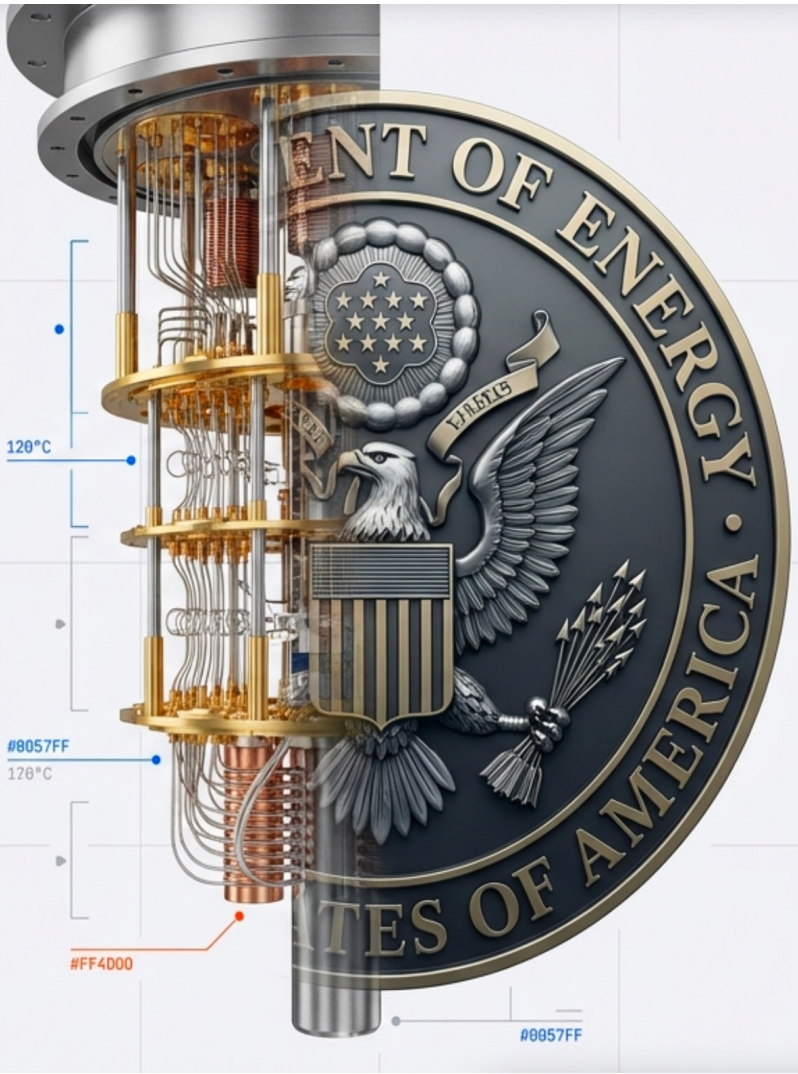

AUTHORITY: Executive Order | November 24, 2025
LEAD AGENCY: Department of Energy (DOE) + White House OSTP
LEADERSHIP: Dr. Darío Gil, Under Secretary for Science

The Third Great Era of American Science

1940s

The Manhattan Project

1960s

The Apollo Program

2025

The Genesis Mission

The Precedent

The Executive Order frames the race for AI dominance as comparable in urgency and ambition to the Manhattan Project and the Apollo Program.

The Mandate

"The largest marshalling of federal scientific resources since the Apollo program."

— Michael Kratsios, Director, White House OSTP.

The Goal

To ensure the United States defines the future of the 21st century by doubling the productivity and impact of American research within a decade.

Breaking the Innovation Plateau

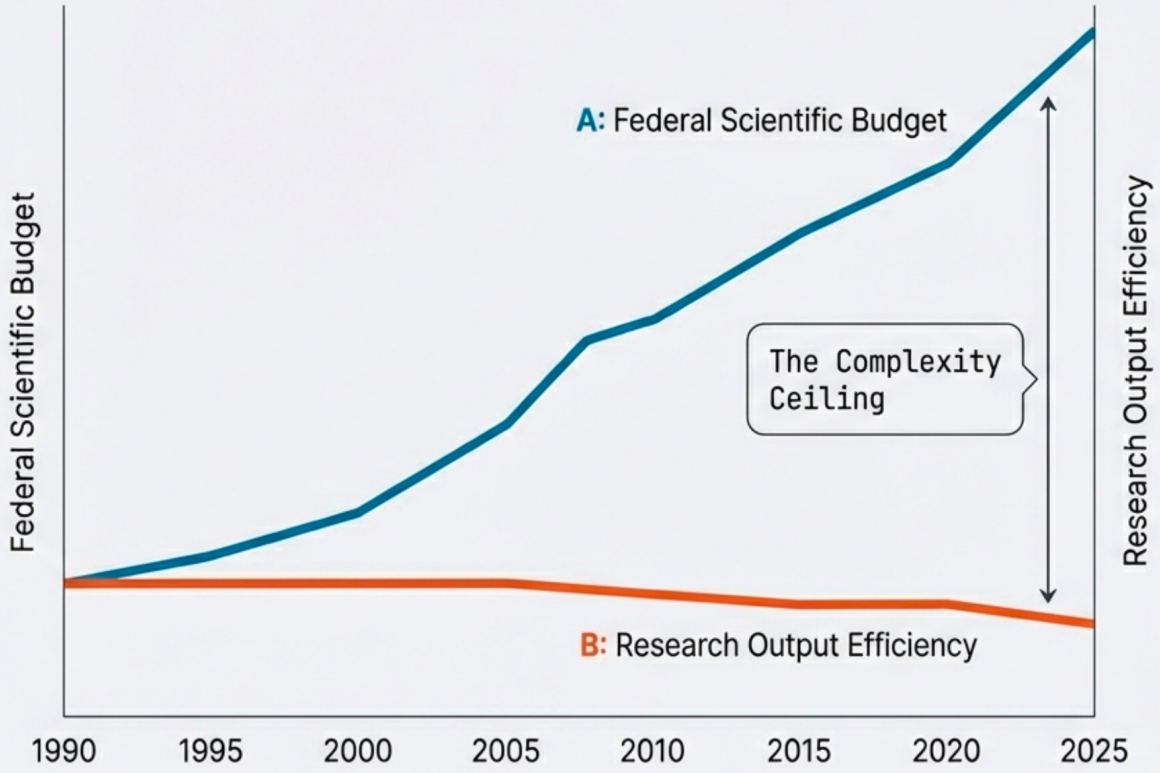

The Context

Since the 1990s, America's scientific edge has faced growing structural challenges. Despite soaring budgets, critical metrics like drug approvals and research outputs have declined.

Modern science has become too complex for human-only processing.

The Pivot

The Genesis Mission is not just about funding; it is about methodology. It seeks to reverse stagnation by unifying agency scientific efforts and integrating AI as a scientific tool.

The Third Great Era of American Science

1940s — 1960s — 2025

doubling the productivity and impact of American research within a decade.

The Manhattan Project

The Precedent

The Executive Order frames the race for AI dominance as comparable in urgency and ambition to the Manhattan Project and the Apollo Program.

The Apollo Program

The Mandate

"The largest marshalling of federal scientific resources since the Apollo program."

— Michael Kratsios, Director, White House OSTP.

The Genesis Mission

The Goal

To ensure the United States defines the future of the 21st century by doubling the productivity and impact of American research within a decade.

PROJECT: AMERICAN SCIENCE INITIATIVE | DOC. REV: 3.1 | DATE: OCT 26, 2023

- **DATA**: American scientific data should be accessible through intuitive, secure and interoperable platforms. We must have the world's largest collection of scientific data from consenting parties, including making existing data AI-ready and building infrastructure to collect new scientific data points at scale.
- **COMPUTE**: Researchers in the U.S. must have access to the world's largest supply of advanced computing, including the supporting networking and energy infrastructure.
- **AI**: We must develop revolutionary AI capabilities to accelerate science, including scientific copilots that innovate routine tasks and leave scientists free to focus on work that requires human ingenuity. The U.S. will set the stage for responsible development and use of artificial general intelligence capabilities that will eventually unlock unparalleled opportunities for the American people.
- **COLLABORATION**: There is immense, untapped brilliance across our nation. We must enable people with expertise from across disciplines to collaborate in advancing American science.
- **PROCESS**: Streamlining regulatory processes and adopting AI-enabled automated testing protocols can result in faster time-to-market for new products while exceeding current safety and efficacy standards. We must build the capacity to get vaccines and therapeutics to market in weeks or months rather than years.

DOE's Genesis Mission is a major new National AI Initiative

- To transform how science is conducted in the US
- Combines the nation's most advanced computing, frontier AI, and world-class scientific instruments
- Targets urgent challenges in energy, discovery science, and national security
- Reimagines the scientific process—automating data-intensive workflows, rapidly generating and testing hypotheses, and enabling discoveries with unprecedented precision and impact
- Integrates HPC, AI, and quantum technologies into a unified innovation ecosystem
- Builds new, faster pathways for partnerships across industry, academia, and national labs

The Strategic Imperative

Why DOE? Why Now?

The Opportunity

Current commercial AI excels at pattern matching and language. The new frontier is **Agentic and Physics-Informed AI**—systems capable of causal reasoning, respecting conservation laws, and navigating multi-scale dynamic systems.

The Unique DOE Asset

We possess the unique convergence required for this leap:

- World-class User Facilities (Equinox, Synchrotrons)
- Massive Proprietary Datasets (80+ years of nuclear data)
- Deep Domain Expertise

The Mission

To guide innovation through the 'Valley of Death'—bridging the gap between scientific discovery and commercially viable deployment in manufacturing, energy, and defense.

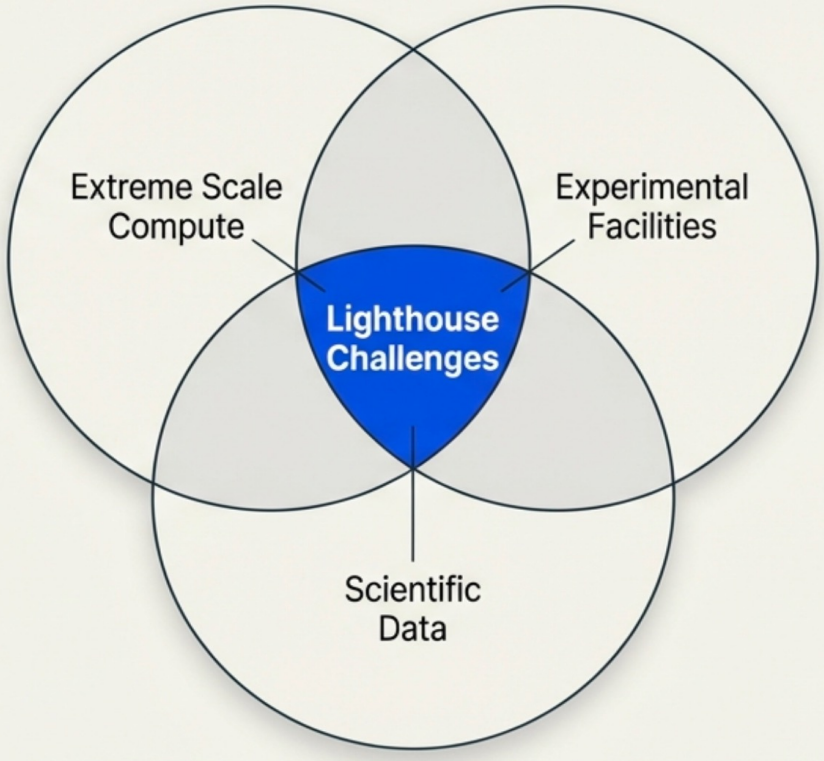

Recent Activities in support of the Genesis Mission

- Funding from the OBBBA (HR1) used to launch
  - "**AI Transformational Model Consortium**" (Lead: Argonne)
  - "**American Science Cloud**" (Lead: ORNL)
- Selection of Initial AI Model Teams to accelerate AI applications across Office of Science and Applied Energy missions (October)
- Formation of integrated teams (HQ/Labs) around AI pillars (Applications, Data, Models, Infrastructure, Partnerships) across SC/AE and NNSA (October)
- National Science and Technology Challenges identified (February)
- RFA to address National Science and Technology Challenges issued March 17
  - **Many** proposals submitted in response

Energy Department Announces 26 Genesis Mission Science and Technology Challenges to Accelerate AI-Enabled American Innovation and Leadership

February 12, 2026

https://www.energy.gov/articles/energy-department-announces-26-genesis-mission-science-and-technology-challenges

Securing America's Critical Minerals Supply

Challenge: America's dependence on foreign supply chains for CMM threatens national security, economic competitiveness, and the deployment of technologies essential for energy independence. Domestic critical mineral production is expensive, complex, and time-consuming, in part because of the many steps to identify, extract, refine, and concentrate from complex, heterogeneous sources across critical mineral supply chains.

AI Solution: AI will revolutionize the entire critical minerals supply chain and development of alternative materials by integrating geophysical data, other fundamental science data, process optimization, cost estimation, and economic modeling into one connected system. Solving this challenge demands an AI that can reason scientifically, can understand complex structure-property relationships, and can design alternatives with different compositions. Physics-based AI offers advanced predictive capabilities to identify alternatives and understand processes underlying critical mineral availability, recovery, refinement, and replacement.

Justification: DOE's existing minerals characterization datasets (e.g., METALLIC, Critical Materials Innovation Hub), combined with DOE national laboratory expertise and DOE-supported efforts in materials science, chemistry, geosciences, biology, process engineering, and economic modeling, could enable acceleration from the years-long mineral development timelines to rapid resource assessment and production optimization. Further, use of AI could reveal new strategies to replace and/or eliminate the need for CMMs in some materials and chemical processes.

National Impact: This effort will reduce reliance on adversarial nations, expand America's mineral resource base, maximize production profitability, and strengthen supply chain resilience for technologies essential to national security and economic prosperity.

Predicting U.S. Water for Energy

Challenge: Water availability is essential for expanding production and utilization of energy, as well as the nation's health and security. However, there are fundamental scientific gaps in our understanding of terrestrial and atmospheric systems that limit our ability to predict water resources, especially on the time scale of weeks to years.

AI Solution: AI capable of multi-scale temporal reasoning could tackle three inter-related grand challenges: cloud physics, surface and subsurface water flows, and the broader hydrologic cycle. AI could improve, accelerate, and couple exascale-class modeling systems through advances in model initialization, and develop surrogates trained on DOE's atmospheric and terrestrial observations and laboratory data, at a fraction of the computational cost of existing models. AI-based model diagnostics for enhanced analysis could refine a model-observational system aligned with decision-making needs.

Justification: DOE is the only agency with AI expertise, advanced computing, and integrated modeling capabilities (e.g., the DOE Energy Exascale Earth System Model, or E3SM), and infrastructure for field research necessary to meet the challenge of providing accurate information on surface and ground-water availability on the time scales of weeks to years.

National Impact: Solutions to these longstanding science challenges will radically improve America's ability to anticipate water supply in the context of changing water availability, demands, energy technologies, and ambitions for energy expansion.

Genesis Mission Teams

INTEGRATED TEAMS

American Science Cloud (AmSC)
Federated infrastructure & computing

Transformational AI Models Consortium (ModCon)
AI models & domain expertise

AI for National Security (AI4NS)
Mission-critical applications

CROSSCUTTING PILLARS

Infrastructure
Exascale computing & facilities

Data
AI-ready pipelines & standards

Models
Domain AI & foundation models

Partnership
Industry, academia, international

Transformational AI Models Consortium (ModCon)

Mission

- Establish a consortium to accelerate the technical development and scientific discovery of the Model Teams
- Develop and deliver domain cross-cutting services as an engine for transformational AI model development
- Convene partners from industry, academia, and internationally to accelerate AI development and adoption

Four Core Teams

deliver support to Model Teams

IPPF
Partnerships and IP Management

DBS
Data Brokers & Standards

BPSW
Best Practices for Scientific Workflows

BASE
Cross-Cutting AI Capabilities

What We Deliver

AI-Ready Data Pipelines
Transform raw scientific data into training-ready datasets

Scientific Workflows
Leaderboards, workflows, and upskilling for continuous discovery

Transformational Capabilities

- Core Agentic Framework
- Self-Improving Models Harness
- Multimodal Reasoning Frontends
- Safety, Security, Assurance
- Evaluation

Deployed on Genesis Mission Infrastructure
The American Science Cloud

Partners
17 DOE Labs + Industry + Academia

Genesis Mission Mapping
Data and Model Teams

Base R&D is Delivering Baseline AI Capabilities to the Genesis Mission

- Providing reusable methods, frameworks, and tools that underpin every Model Team
- Responsive to the requirements set by the Genesis Mission and ModCon Core teams
- Collaborating directly with Model Teams and AmSC

We are also:

- Genesis Mission's crosscutting AI dev team
- Filling in what is missing from the current AI ecosystem
- Helping inspire others to use the command line

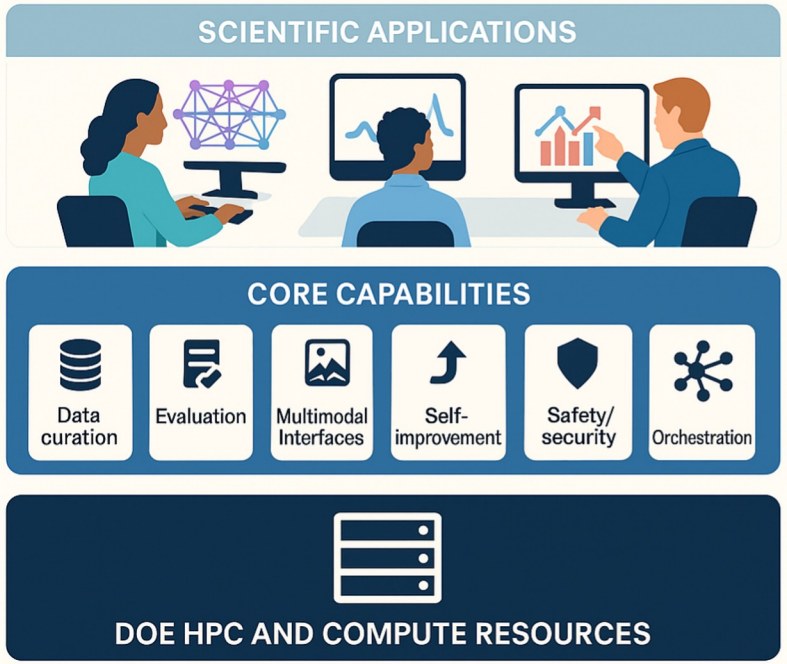

FASTER SCIENCE

SAFER AI

SCALABLE ORCHESTRATION

Best Practices for Scientific Workflows
BPSW Team

Define **requirements** for cross-domain agentic frameworks and network topologies, in coordination w AmSC and BASE.

Rigorously **evaluate** then optimize how AI methods accelerate DOE science tasks and full workflows.

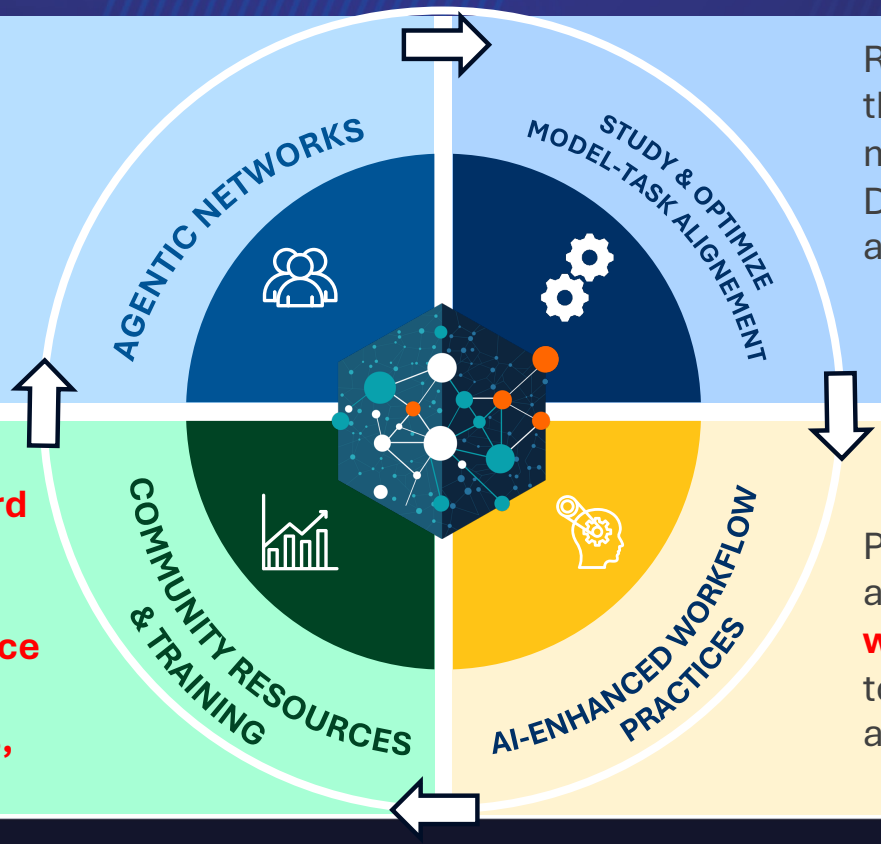

Maintain a **leaderboard** ranking AI models. A challenge repository spanning, **best-practice guides, tutorials, AI integration examples, and success stories.**

Publish **guidelines** and reusable **agentic workflow templates** to standardize AI adoption across DOE

Data Brokers and Standards
DBS Team

Standards
Define data standards and access requirements for ModCon and Model Teams in coordination with AmSC

Raw to AI Ready (RA→RAI) Pipeline Specs
Design and specify the raw to AI-ready data pipeline (RA→AIR) using model team driven requirements.

Network of Expertise
Support Model Teams through a broad network of experts who understand model training & fine-tuning, alongside Data Detectives

Evaluation
Evaluate data-related proposals from data generators and Model Teams.

Private Public Partnerships (P3) Team

16 Model Teams using transformational AI capabilities to address national challenges

14 INITIAL MODEL TEAMS

AIForHPC
AI-Driven Code Dev | ANL

Microelectronics
AI Co-Design | ANL

AI4Quantum
Quantum Algorithms | ANL

COMB-FLOW
Combustion & Fluids | ANL

CM2US
Critical Minerals | ANL

Quarks to Cosmos
HEP & Cosmology | ANL

AXESS
Specs-to-Silicon | FNAL

Prometheus
Nuclear Power | INL

SYNAPS-I
Light Sources | LBNL

MOAT
Accelerators | LBNL

GridAI
Power Grid | NREL

MAIQMag
2D Quantum Magnets | ORNL

FAMOUS
Biological Systems | PNNL

Fusion-FM
Magnetic Fusion | PPPL

2 TEAMS Integration and 1 expanding

ISAAC
Scientific User Facility | SLAC

Alphafold for Microelectrontics
BES

MAIQMag expansion to include QCAD
Quantum Co-Design

The American Science Cloud (AmSC)

- Led by ORNL with Argonne and other DOE labs as partners.

- Goals:
  - Develop the software and hardware infrastructure for DOE's AI data and model development efforts
  - Develop an API that abstracts away the complexity of accessing data and executing computational and experimental workflows

- ModCon will work closely with AmSC. DOE AI models will be provided to the scientific community through AmSC.

AmSC
AMERICAN SCIENCE CLOUD

New AI Systems coming to Argonne

- Partnership between Argonne, Nvidia, Oracle
- **Equinox**: A system with **10,000** NVIDIA Blackwell GPUs to be deployed soon
- **Solstice**: A system with **100,000** NVIDIA Blackwell GPUs scheduled to be deployed in 2027
- Solstice will be the largest AI supercomputer in the DOE complex
- https://www.energy.gov/articles/energy-department-announces-new-partnership-nvidia-and-oracle-build-largest-doe-ai

And a few other Systems

- **Minerva**: For ALCF AI inference service; 64 NVIDIA Blackwell GPUs, NVIDIA Quantum-2 InfiniBand interconnect
- **Janus**: For development of the next-generation AI and HPC workforce; 64 NVIDIA Hopper GPUs and BlueField-3 DPUs
- **Tara**: For large-scale AI inference; 2,688 NVIDIA GH200 Grace Hopper GPUs

Trillion Parameter Consortium

TPC
Trillion Parameter Consortium

TPC Goals

- ***Build an open community*** of researchers creating state-of-the-art large-scale generative AI models for science and engineering,
- ***Incubate, launch, and facilitate*** coordination and collaboration for specific projects building such models, and
- ***Create a global network*** of resources and expertise to facilitate teaming and training of next-generation AI researchers.

TPC Operating Principles

- Transparency
- Fairness, and
- Ethical AI practices, including

TPC goals and principles align with scientific and government guidelines, promoting transparency, mitigating bias, ensuring trustworthiness, protecting privacy, fostering collaboration, and embracing adaptability in our AI development.

TPC26 Conference

https://tpc26.org

3rd Conference on Foundation Models and AI Agents for Science

David Rubenstein Forum, 1201 E 60th Street, **Chicago, IL**

May 27-29, 2026

SciFM 2026 is organized around three interlocking questions

1. Can foundation models be grounded in physical law — and what does 'grounded' actually mean for a scientific model?
2. What are the specific, unsolved problems in aerospace, energy, and manufacturing that a physics-aware AI could address — and why hasn't the AI industry addressed them yet?
3. What does it look like when AI participates in scientific discovery at facility scale, in real time, with appropriate uncertainty and human oversight?

https://www.scifmconferences.org/

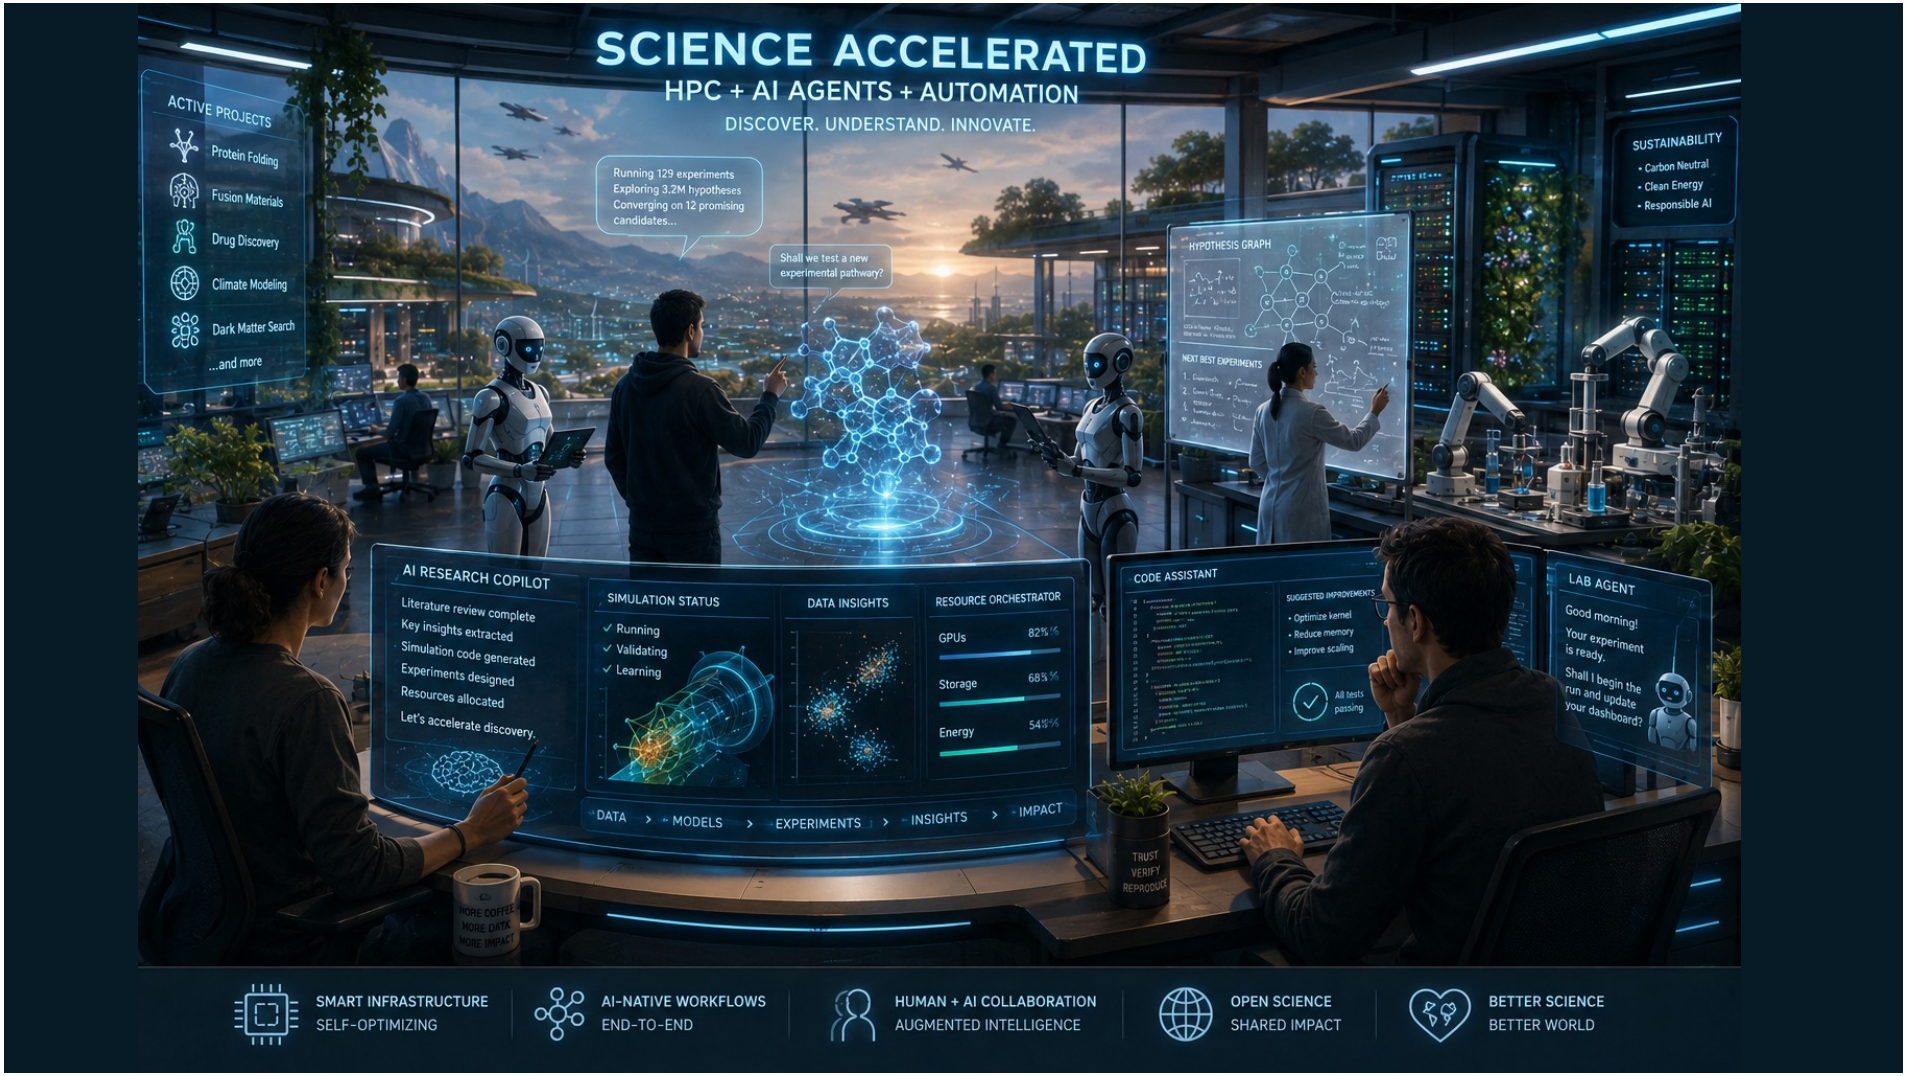
SCIENCE ACCELERATED
HPC + AI AGENTS + AUTOMATION
DISCOVER. UNDERSTAND. INNOVATE.
ACTIVE PROJECTS
Protein Folding
Fusion Materials
Drug Discovery
Climate Modeling
Dark Matter Search
...and more
Running 129 experiments
Exploring 3.2M hypotheses
Converging on 12 promising candidates...
Shall we test a new experimental pathway?
HYPOTHESIS GRAPH
NEXT BEST EXPERIMENTS
SUSTAINABILITY
• Carbon Neutral
• Clean Energy
• Responsible AI
AI RESEARCH COPILOT
Literature review complete
Key insights extracted
Simulation code generated
Experiments designed
Resources allocated
Let's accelerate discovery.
SIMULATION STATUS
✓ Running
✓ Validating
✓ Learning
DATA INSIGHTS
RESOURCE ORCHESTRATOR
GPUs 82%
Storage 68%
Energy 54%
CODE ASSISTANT
SUGGESTED IMPROVEMENTS
• Optimize kernel
• Reduce memory
• Improve scaling
All tests passing
LAB AGENT
Good morning!
Your experiment is ready.
Shall I begin the run and update your dashboard?
DATA > MODELS > EXPERIMENTS > INSIGHTS > IMPACT
MORE COFFEE
MORE DATA
MORE IMPACT
TRUST
VERIFY
REPRODUCE
SMART INFRASTRUCTURE
SELF-OPTIMIZING
AI-NATIVE WORKFLOWS
END-TO-END
HUMAN + AI COLLABORATION
AUGMENTED INTELLIGENCE
OPEN SCIENCE
SHARED IMPACT
BETTER SCIENCE
BETTER WORLD