

Accelerated Knowledge Discovery: A Vision for NASA Science

Rahul Ramachandran, Kaylin Bugbee and Juan Bernabe-Moreno

NASA/MSFC
IBM Research Europe

Key Points

- Accelerated Knowledge Discovery (AKD) uses AI as a partner in scientific research to speed up the entire discovery process
- AKD enables continuous research cycles by automating tasks like data analysis, hypothesis generation, and reporting, reducing bottlenecks
- AKD can help solve complex scientific problems faster by combining AI with existing data, tools, and research goals

Abstract

This paper introduces the sixth paradigm of scientific discovery: accelerated knowledge discovery (AKD). This paradigm is defined by the full integration of AI into the research workflow as a tool augmenting cognitive capabilities of human scientists. AKD emerges from the convergence of advanced AI models, autonomous agentic systems, and human-AI collaboration.

AKD accelerates the research cycle by reducing the time from conceptualization to discovery. It automates labor-intensive tasks such as literature review, hypothesis generation, experimental design, data analysis, modeling, simulation, and manuscript drafting. In well-defined domains, AKD can transform the scientific method into a continuously adaptive cycle, where outputs from each phase inform the next. These closed-loop scientific workflows shorten discovery timelines and reduce overhead. In addition to the scientific speed up, AKD targets an increase in the quality of research allowing for more systematic discovery of knowledge.

However, AKD's success depends on principled, trustworthy design. This requires a holistic approach that emphasizes explainability, reproducibility, robustness, adaptability, and transparency. Key requirements include alignment with open science principles, required human oversight, scientific accountability, and rigorous provenance tracking. Human researchers must remain ultimately responsible for scientific integrity, ethical reasoning, and interpretation, with AI serving as an augmentative partner.

Although the proposed approach applies to any domain focused on scientific discovery, this paper highlights AKD's potential to advance NASA's science mission, given the agency's vast data assets, complex objectives, and interdisciplinary challenges. By integrating NASA's data, foundation models, scalable computing, and knowledge frameworks, AKD can accelerate discovery and foster innovation across its scientific portfolio.

Plain Language Summary

This paper introduces a new way of doing science called *accelerated knowledge discovery* (AKD). AKD uses artificial intelligence (AI) as a tool to significantly augment cognitive capabilities of scientists to help make discoveries faster. With recent advances in AI, like powerful language models and smart software agents, scientific research can become more efficient and more creative.

AKD helps speed up many parts of the research process, including reviewing past studies, coming up with new ideas, designing experiments, analyzing data, building models, and even writing up results. In some cases, these AI systems can run a full cycle of scientific work almost automatically, where each step helps improve the next, reducing time and effort while opening the door to new kinds of questions and discoveries.

While AKD can be applied to any scientific domain, this paper explores how it could be especially useful for NASA space science, which deals with massive amounts of data and complex scientific challenges. By combining AI with NASA's data and research tools, AKD could help solve big problems faster and support breakthroughs across many areas of science.

Introduction

The scientific process has continually evolved in response to the emergence of new tools, technologies, and data sources. These evolutions have been driven by paradigm shifts in scientific methodologies, each one expanding the boundaries of inquiry and redefining what is knowable. From early empirical observations to the modern frontier of AI-assisted research, each paradigm has enabled scientists to ask fundamentally new types of scientific questions. Understanding this progression is essential to anticipating and shaping the next transformative change.

We are now approaching another inflection point: the full integration of AI into the scientific research workflow. This position paper describes the impending shift and argues that, to be successful, the transformative technologies in development must be grounded in principled, trustworthy design and aligned with responsible AI and open science practices.

AI Integration into Scientific Processes

Recent research demonstrates that AI, especially large language models (LLMs), is already transforming key steps of the scientific process. For instance, Ren et al. (2025) showed that LLM-powered scientific agents can employ prompt-based planning, utilize multiple memory modalities (e.g., historical context, external knowledge bases), and interact with toolsets to perform complex tasks such as literature review, hypothesis generation, and information synthesis. These agents streamline labor-intensive processes like information retrieval, screening, and summarization, enhancing the reasoning capacity of standalone LLMs and infusing efficiency into research workflows.

Another study by Dawid et al. (2025) further highlights how domain-specific AI agents, when combined with appropriate human oversight, can accelerate many aspects of scientific inquiry. Importantly, these systems are designed to augment—not replace—human researchers. They aim to enhance scientific productivity while maintaining rigor and accountability. For example, AI Scientist-v2 (Yamada et al., 2025), an advanced agentic system, demonstrated the ability to autonomously generate a peer-reviewed workshop paper. While this system exhibited high autonomy, human researchers were still responsible for selecting the best ideas and final output, illustrating a collaborative approach guided by human insight. Similarly, the AI Cosmologist (Moss, 2025) has automated workflows in astronomy and cosmology research by integrating agents for literature search, data analysis, and result synthesis. The system can generate LaTeX-ready manuscripts, visualizations, and bibliographies, reducing the time and effort required to communicate research results in complex domains.

The Sixth Paradigm: Accelerated Knowledge Discovery

Inspired by these advancements, we introduce the sixth paradigm, accelerated knowledge discovery (AKD), characterized by the convergence of advanced AI models, autonomous agentic systems, and human-AI collaboration. AKD transcends previous paradigms by integrating all modes of inquiry into a coherent, intelligent, and iterative scientific process.

Historically, scientific progress can be characterized as evolving via five paradigms: (1) empirical observation, (2) theoretical modeling, (3) computational simulation, (4) data-intensive discovery, and (5) AI/ML-driven inference (Gray, 2007; Yannis, n.d.). The fourth paradigm emphasized managing massive datasets and introduced scalable data infrastructure, metadata standards, and the prioritization of reproducibility (Gray, 2007). The fifth paradigm built on this evolution by focusing on the use of machine learning to develop predictive models and identify complex patterns, and enabling innovations like digital twins (Yannis, n.d.).

The emergence of AKD as the sixth paradigm signals a shift in how AI is and will be used, not just as an analytical engine, but as a cognitive tool to augment human thinking. Agentic systems now exhibit contextual awareness, decision-making capability, and the ability to execute workflows across the entire scientific lifecycle, from hypothesis formulation and experimental design, to data interpretation and knowledge refinement. Unlike earlier paradigms, AKD integrates all previous modes of inquiry into a cohesive, intelligent, iterative process.

This reframing positions AI as a “thinking tool,” augmenting human insight and intuition with computational power and automation. Human researchers guide scope, define scientific constraints, and bring ethical reasoning and creativity, while AI agents contribute speed, scalability, and reasoning support. In this collaborative paradigm, AI operates as a cognitive enhancing tool working across data-rich and tool-diverse ecosystems such as those found at NASA.

This paper outlines the concept of Accelerated Knowledge Discovery anchored in responsible AI and open science. The paper also discusses how AKD can support NASA’s science goals by supporting complex research workflows and leveraging the existing data and information systems.

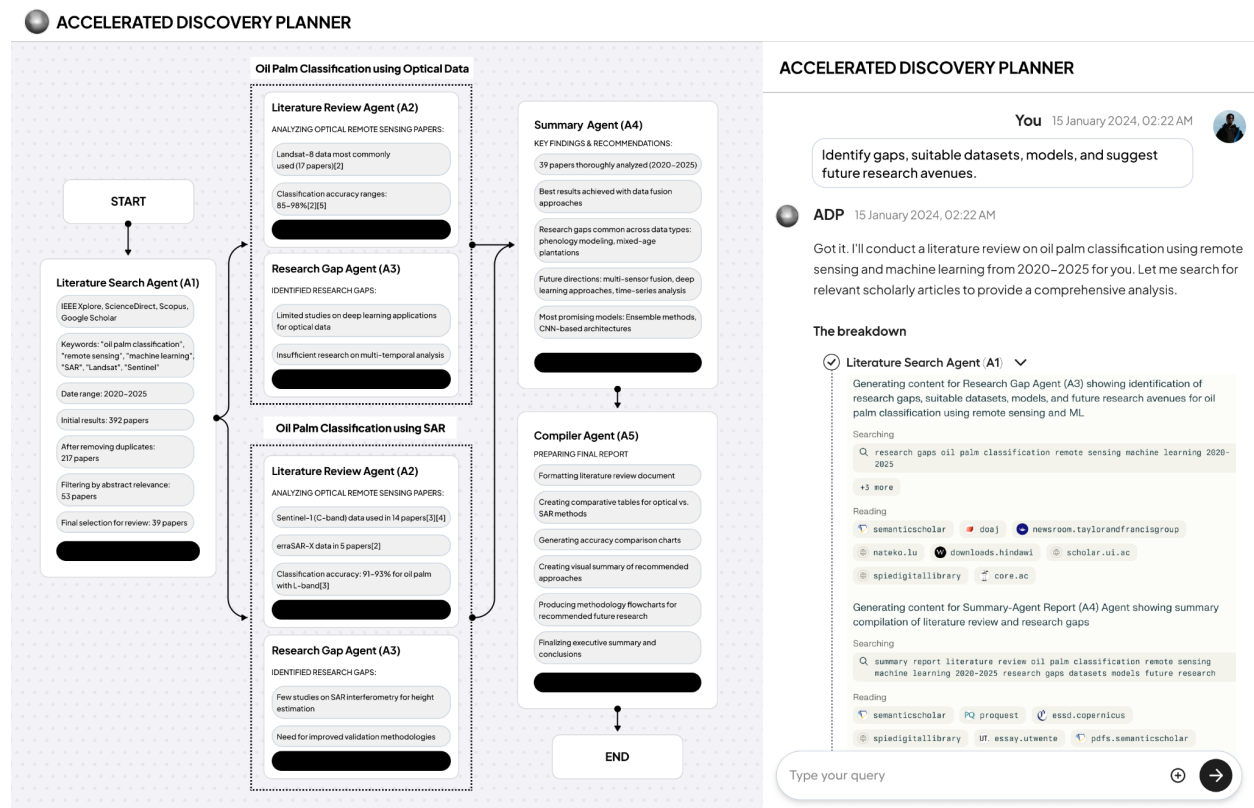


Figure 1. Illustrative mockup of an accelerated knowledge discovery (AKD) system in use. In this scenario, a researcher is investigating oil palm classification and engages the AKD system to identify current knowledge gaps, recommend suitable datasets and models, and suggest new avenues for future research. This example highlights how AKD can intelligently orchestrate the research process by synthesizing existing knowledge, surfacing relevant resources, and proposing actionable next steps to expedite scientific discovery. Researchers have control over

all the steps in this workflow to enable Human in the Loop (HITL). Each step can be evaluated and modified.

Science at NASA

NASA's science enterprise is vast, covering a broad spectrum of research areas from understanding Earth's complex systems to exploring distant galaxies and investigating the fundamental origins of life. While exploration remains a core part of its identity, NASA's scientific scope ("NASA's Science Vision," n.d.) extends well beyond discovery for its own sake. NASA's science vision emphasizes a commitment to generating actionable knowledge that informs public policy, supports environmental stewardship, advances technological innovation, and deepens our understanding of the universe. As such, NASA's mission inherently addresses fundamental scientific questions about our planet, our solar system, and the broader cosmos. These central questions, shaped by evolving scientific priorities and global challenges, inspire NASA's science ventures: How can we monitor and understand the complexity of Earth systems? How does the Sun shape the heliosphere and influence planetary environments? How do planets and the building blocks of life emerge from cosmic origins? What can space-based biological and physical experiments reveal about life's fundamental processes? And, ultimately, are we alone in the universe?

Seeking answers to these profound questions requires coordinated, interdisciplinary research that integrates observational science, theoretical modeling, high-performance simulation, and rigorous experimentation across a wide range of domains. NASA's Science Mission Directorate (SMD) is responsible for managing this complex research portfolio. The SMD oversees mission planning, fosters collaboration, and supports scientific leadership across five primary divisions: Astrophysics, Biological and Physical Science, Earth Science, Heliophysics, and Planetary Science. These divisions rely on a diverse and sophisticated array of platforms including space telescopes, planetary rovers, Earth-observing satellites, and the International Space Station (ISS) to collect critical scientific data.

The data volume generated by these platforms is immense and perpetually increasing. Over time, even a single mission can generate petabytes of data consisting of navigation measurements, multimodal and multi-scale observations, as well as calibration and validation datasets. Managing this level of data complexity demands robust infrastructure for end-to-end data stewardship that includes meticulous curation and development of data discovery, access, and analysis solutions. In alignment with its Open Science mandate, NASA ensures that its extensive data holdings are publicly available via platforms such as Earthdata ("Your Gateway," n.d.), the Heliophysics Data Portal ("Heliophysics Data Portal," n.d.), the Planetary Data Ecosystem (PDE) ("Planetary Data Ecosystem," n.d.), Open Science Data Repository (OSDR) ("NASA Open Science Data Repository," n.d.), NASA/IPAC Infrared Science Archive ("IRSA," n.d.) and the cross-disciplinary Science Discovery Engine (SDE) ("Science Discovery Engine," n.d.).

However, the scale, heterogeneity, and disciplinary fragmentation of NASA's scientific data ecosystem present formidable challenges. Extracting insights from these rich datasets requires more than conventional analytical approaches. This is precisely the environment where agentic AI systems, such as those envisioned under the AKD paradigm, offer transformative potential. These intelligent systems can orchestrate complex scientific workflows, coordinating across disparate data sources, computational tools, and disciplinary models, to spur discovery, improve efficiency, and support NASA's multifaceted science mission.

Accelerated Knowledge Discovery (AKD)

The emergence of the sixth paradigm of AKD is not the result of a single technological breakthrough. Rather, the new paradigm reflects the convergence of multiple, maturing technologies that span artificial intelligence (AI), modern data infrastructures, and evolving scientific workflows. These technologies have progressed in parallel, creating the conditions necessary to build intelligent, connected, and highly adaptive systems. While the fifth paradigm was marked by the adoption of AI tools for data-driven modeling and analysis, AKD employs AI in a fundamentally different way. Research processes infused with AKD leverage generative and reasoning-based AI as tools that can augment human thinking and help co-design scientific workflows. This role shift from any passive tool to active tool to enhance cognition capabilities in research is a defining characteristic of AKD.

AKD is inherently integrative, bringing together advanced AI models (e.g., large language models and planning agents), scalable computational infrastructure (e.g., cloud computing), and structured knowledge frameworks (e.g., knowledge graphs). This technological convergence enables the development of scientific workflows that are context-aware, iterative, and capable of supporting sophisticated reasoning across disciplines. AKD systems are not just faster, they are more adaptive, enabling a fundamentally new approaches to knowledge discovery.

For an organization like NASA, the significance of AKD is considerable. NASA's broad spectrum of missions, the diversity of its observational instruments, and the volume and complexity of its data ecosystem demand more than piecemeal solutions. AKD provides support with crafting experiments, analyzing multimodal datasets, and bridging disciplinary boundaries by strategically employing technology. These agentic systems are capable of navigating NASA's vast, heterogeneous data and tool landscape, revealing scalable, responsive pathways to insight.

Several converging technologies underpin AKD's capabilities. First, generative AI and LLMs function as powerful reasoning engines that can interpret complex natural language prompts and synthesize knowledge from large, diverse sources. Second, domain-specialized foundation models, trained on NASA's multimodal scientific data, enable rapid deployment of tailored applications such as real-time flood detection or predictive solar flare forecasting. Third, agentic AI frameworks support intelligent agents capable of decomposing research goals into sub-tasks,

planning workflows, and autonomously executing scientific tasks using available tools and services.

Another foundational component of AKD is explicit knowledge representation. Knowledge graphs (KGs) formally describe entities, relationships, and properties across scientific domains, forming a structured layer for AI reasoning. These representations enable AI systems to detect patterns, reason across heterogeneous data sources, and integrate diverse information streams while adhering to scientific constraints. In effect, KGs act as both map and filter, providing the contextual grounding needed for scientifically coherent automation.

To fully realize AKD's potential, a robust computational infrastructure is also essential. High-performance computing (HPC) resources will be required for developing, training, and fine-tuning large-scale AI models for specific scientific domains. In parallel, cloud platforms will play a critical role in enabling flexible, on-demand inferencing and scalable deployment of agentic workflows. This combination of HPC and cloud computing forms the backbone of AKD, providing the adaptability and throughput needed to support real-time, data-intensive, and scalable scientific processes.

Ultimately, AKD represents a shift away from tool-centric workflows toward integrated, intelligent systems that learn, reason, and iterate in collaboration with human scientists. The paradigm offers a systematic approach for transforming the speed, scale, and reproducibility of scientific discovery, making AKD particularly well-suited to the complexity of NASA's mission-driven research challenges.

Reimagining Scientific Workflows with AKD

The sixth paradigm of accelerated knowledge discovery is defined by its symbiotic human-AI design, using AI to augment human intellect and insight rather than replacing it. AKD blends AI's speed and scale with uniquely human creativity, judgment, and contextual awareness, making human-in-the-loop (HITL) design a foundational principle. Human researchers guide inquiry, provide oversight, and interpret results, while AI agents manage computations, data analysis, and workflows. This approach transforms scientific workflows into adaptive, iterative research cycles. In AKD, AI remains an assistant to human expertise, not a replacement.

Throughout the scientific lifecycle, AKD enables a reimagined partnership between researchers and intelligent systems. During the proposal development phase, AI agents can help formulate competitive grant applications, including generating required open science data management plans (OSDMPs) and structuring proposals for observational campaigns, such as those requesting time on major astronomical platforms like the Hubble or James Webb Space Telescopes. These agents can suggest appropriate funding opportunities, ensure alignment

with agency guidelines, and assist with administrative formatting, allowing scientists to focus on articulating their core research objectives.

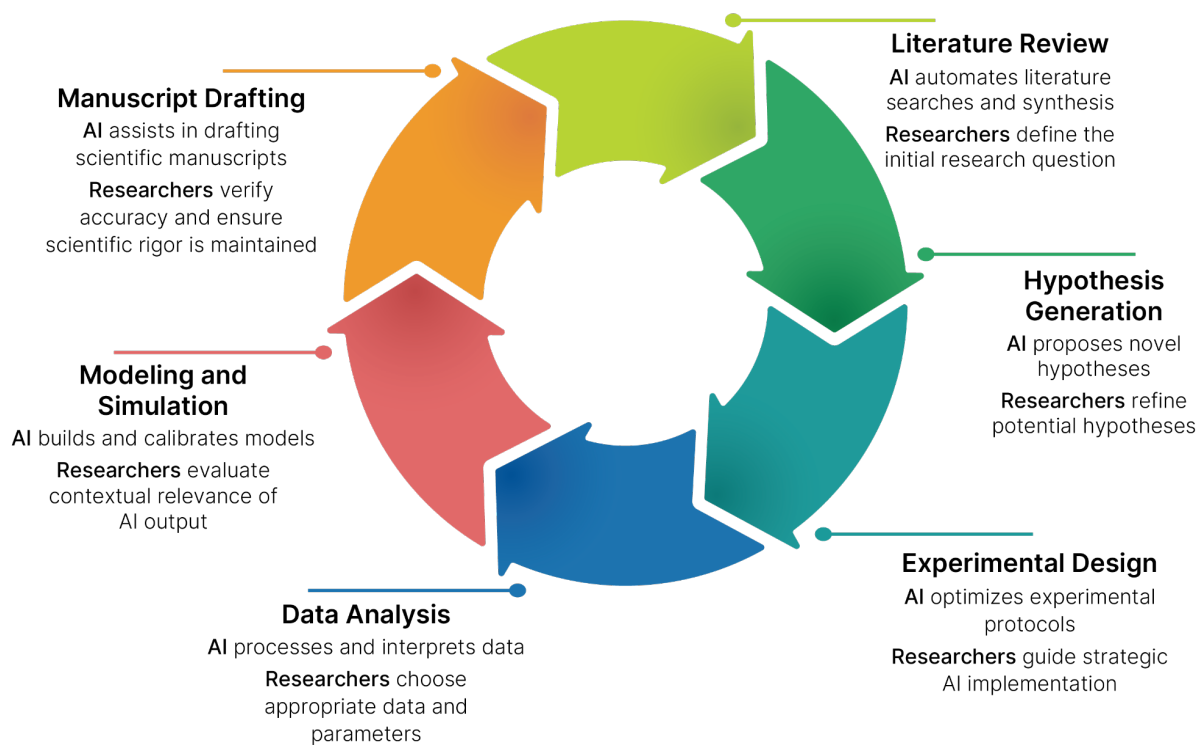


Figure 2. Illustration of the scientific discovery process and key points of AI integration. This figure depicts the major stages of the scientific research lifecycle, ranging from hypothesis generation to experimentation, analysis, interpretation, and dissemination. The illustration highlights where AI, particularly within the AKD framework, can be fully integrated to augment and accelerate each phase of the workflow.

In the domain of literature review and knowledge synthesis, LLM-powered agents can conduct large-scale searches, extract key findings, identify knowledge gaps, and distill emerging trends from vast scientific corpora. These agents may analyze user-supplied bibliographic databases (e.g., BibTeX files), highlight important excerpts, and generate accessible summaries of complex technical material. However, it is the human researchers who define the research question, guide the AI's strategy, and evaluate the contextual relevance and factual accuracy of the synthesized insights, ensuring scientific rigor is maintained.

For hypothesis generation, AKD agents can identify statistically significant correlations, extrapolate trends, and reason across domains to propose novel, testable hypotheses. When

integrated with structured knowledge graphs, these agents can anchor generated hypotheses in existing scientific understanding, reducing the risk of spurious or implausible suggestions. Scientists, in turn, engage in evaluating the novelty, feasibility, and theoretical soundness of these hypotheses, refine them based on their domain-specific insights, and prioritize those that are most promising for follow-up investigation.

In the experimental design phase, AKD systems can optimize parameters, suggest alternative configurations, and even autonomously design entire experimental workflows in closed-loop systems. They can simulate multiple design scenarios to assess trade-offs and highlight the most informative pathways forward. Human researchers define the experimental constraints, validate AI-generated designs, and adjust for real-world conditions that may not be fully represented in models or simulations.

During data analysis and interpretation, AKD workflows streamline data preprocessing, feature engineering, and modeling, reducing the manual burden on scientists and accelerating time-to-insight. AI agents assist with pattern recognition, anomaly detection, and statistical validation, while human scientists choose appropriate methods, evaluate data quality, interpret the significance of results, and place findings within the broader scientific context.

In the realm of modeling and simulation, AKD systems support the development of data-driven models and the construction of fast emulators for computationally expensive physical simulations. These agents can learn novel functional mappings from empirical or simulated data and aid in the fine-tuning of foundation models (FMs) for specific domains. They also automate the exploration of high-dimensional parameter spaces, enabling comprehensive simulation explorations. Yet, the validation of these models, including their realism, theoretical underpinnings, and boundary conditions, remains in the hands of human experts.

Crucially, AKD must be aligned with open science principles to promote transparency, reproducibility, and collaboration. AI agents can assist in automating documentation, managing metadata, enforcing FAIR (Findable, Accessible, Interoperable, and Reusable) practices, and ensuring that research artifacts such as data, code, models, and workflows are well-curated and openly available. This alignment with institutional and governmental mandates, such as NASA's SPD-41a directive, helps embed openness into the scientific workflow from the beginning.

In the scientific writing and dissemination phase, AI tools within the AKD framework can support researchers by drafting preliminary text, generating figures and tables from data, formatting references, and checking grammar and style. However, the construction of the scientific argument, the articulation of insights, and the responsibility for accuracy and interpretation remain firmly with the human authors. AI serves to streamline the process, not to substitute scholarly authorship.

AKD can also play a role in enhancing peer review. AI agents can assist reviewers by flagging methodological issues, checking citation completeness, evaluating dataset quality, and suggesting relevant literature. Nonetheless, final judgments regarding novelty, methodological soundness, and broader impact continue to rest with human reviewers. Finally, AKD supports

scientist training and continuous learning. AI systems can help researchers better understand scientific papers, technical reports, or codebases by offering guided explanations and tutorials tailored to their existing knowledge. They can also recommend new skill development pathways, such as learning programming languages or statistical techniques, and provide real-time, interactive support for troubleshooting or building conceptual understanding.

Taken together, these capabilities point to the viability of closed-loop scientific workflows (CLSWs) under the AKD paradigm. In such workflows, outputs from one phase, such as automated data analysis, immediately inform the next, such as hypothesis refinement or targeted experimentation. Because these workflows operate within an explicitly defined scope and a curated catalog of tools and services, they can be executed safely, reproducibly, and cost-effectively. Routine, data-heavy tasks such as ingestion, preprocessing, model selection, and metric evaluation can be handled autonomously, while human scientists remain engaged in setting the research agenda, interpreting results, and ensuring that scientific rigor is maintained. Through this tight coupling of automation and oversight, AKD transforms the sixth paradigm from just a concept into an operational reality available today.

| AKD Technology | Purpose | Functionality |
|--|---|---|
| Generative AI & Large Language Models | Synthesize knowledge from vast, diverse data; generates hypotheses, literature reviews, and summaries | Act as the cognitive engine interpreting complex inputs and outputs within workflows |
| Domain-Specific Foundation Models | Tailored AI models specialized in specific scientific domains; facilitate tasks like prediction and anomaly detection | Provides domain specific models and can be , coupled with generative models to accomplish different tasks |
| Agentic AI Frameworks | Intelligent systems that autonomously plan, execute, and manage scientific workflows and tasks | Coordinate the execution of tasks, integrating generative AI outputs and foundation model predictions |
| Explicit Knowledge Representation | Structured knowledge graphs providing formalized descriptions of entities, relationships, and scientific constraints | Provides contextual grounding, guiding agentic decision making and ensuring coherent AI reasoning |

Robust Computational Infrastructure

High-performance computing and scalable cloud platforms enabling computationally intensive tasks

Supports the entire technological ecosystem, ensuring scalability, responsiveness, and real-time capabilities

Table 1. Descriptions of AKD technological components, their functionality, and purpose.

A key application of AKD is the development of closed-loop scientific workflows (CLSWs) which are self-updating research cycles that operate within clearly defined constraints. These workflows are characterized by four essential components. First, they are grounded in a well-defined scientific scope, captured as machine-readable descriptions of research questions, hypotheses, metrics, and accepted data domains. Second, they operate with a limited inventory of tools and services including datasets, models, simulators, and compute functions, each documented with metadata and accessible via stable APIs or function-calling schemas. Third, they include an autonomous engine that adapts static workflows into dynamic processes, refining hypotheses and reallocating computational effort based on evolving insights. Finally, human-in-the-loop safeguards ensure scientific validity, ethical compliance, and interpretability. Within NASA’s operational context, CLSWs powered by AKD can significantly shorten discovery timelines and reduce the overhead associated with manual iterative tasks.

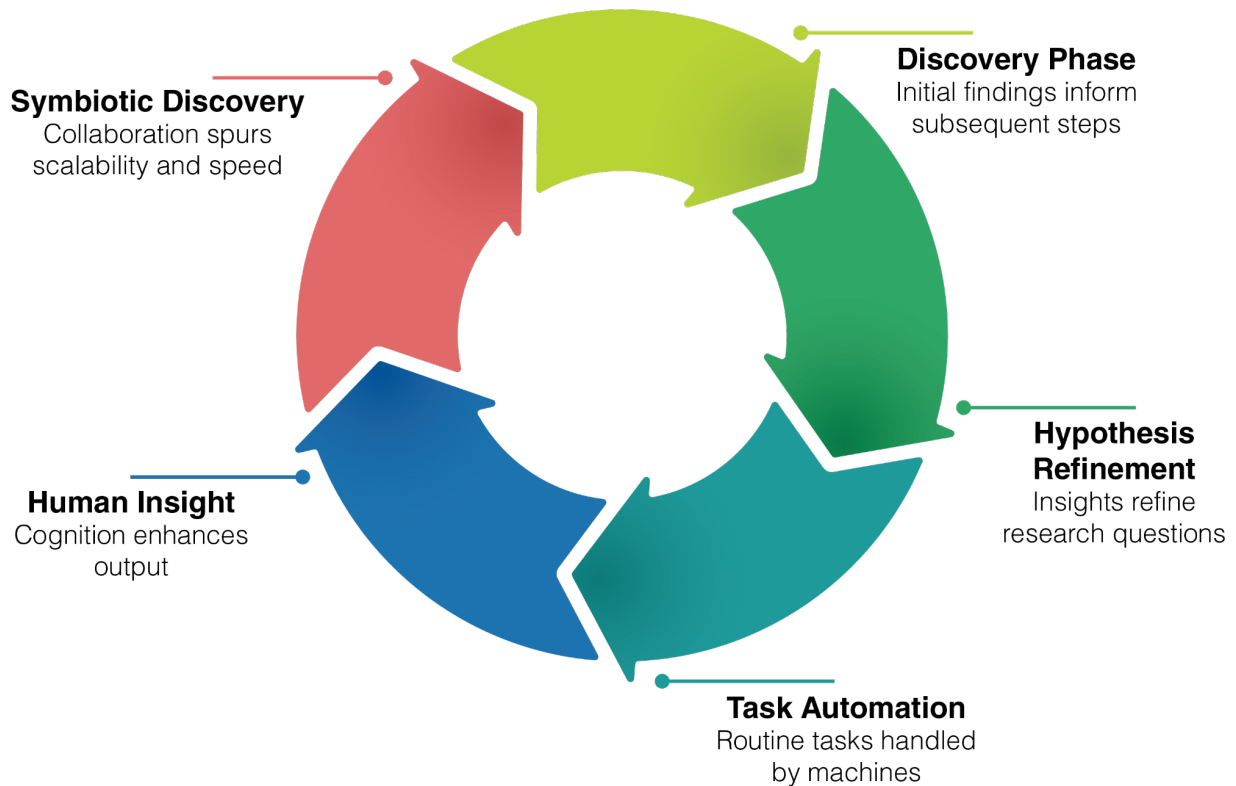


Figure 3. *Conceptual representation of closed-loop scientific workflows (CLSWs). This figure illustrates CLSWs as automated, self-updating cycles of experimentation that operate within clearly defined scientific scopes and tool boundaries. These workflows iteratively refine hypotheses, trigger targeted analyses or simulations, and integrate new data, enabling adaptive and efficient scientific discovery within the AKD framework.*

Building Trust and Mitigating Risks in AKD

The integration of agentic AI into scientific workflows under the AKD paradigm holds tremendous promise for accelerating discovery, but this approach also introduces significant risks that must be proactively mitigated. Without strong ethical foundations and robust operational safeguards, AKD systems may produce flawed or irreproducible scientific results, misallocate research resources, or erode public trust in the scientific enterprise. In the context of tAKD, trustworthiness must encompass far more than model accuracy. Guardrails for AKD framework requires a holistic approach grounded in principles such as explainability, reproducibility, robustness, adaptability, and transparency. These principles are essential to ensuring that AI-generated outputs are both scientifically valid and ethically sound.

While agentic AI systems offer powerful capabilities such as scale, speed, adaptiveness, and even a form of computational creativity, these same strengths can create and exacerbate risks. These include the generation of plausible but incorrect results, automation bias, reduced scientific accountability, and a loss of experimental and computational reproducibility. Thus, realizing AKD's potential requires more than technical advancement; successful implementation demands the development of well-defined, enforceable safeguards. These safeguards must include risk-aware system design, strong human-in-the-loop oversight, and alignment with long-standing norms of scientific integrity. Crucially, the lessons from ongoing discussions in AI ethics must be translated into concrete operational practices embedded directly within the design and deployment of AKD systems.

Core Requirements for Trustworthy AKD

Alignment with Open Science Principles

Open science principles, originally developed to enable collaboration and transparency among human researchers, are even more vital in a future where AI becomes a full participant in the scientific process. Agentic systems developed under AKD must be intentionally aligned with these principles to ensure that the research they support is transparent, interpretable, and verifiable. This alignment not only reinforces trust in scientific findings but also helps prevent the spread of misinformation or low-quality results.

To uphold open science principles (Ramachandran et al., 2021; NASA Open Science Training Team, 2025), AKD systems must provide built-in support for many scientific products:

- **Open Data:** All datasets generated or used by AKD agents must be made openly accessible, reusable, and accompanied by appropriate metadata, attribution, and licensing terms.
- **Open Code:** All software, scripts, and workflow logic developed or executed by AKD systems must be shared under clear open licenses, allowing others to inspect, adapt, or reuse them.
- **Open Results:** Scientific outputs, including intermediate results, workflows, protocols, and technical notes, must be made publicly available in accessible repositories, extending beyond final published papers.
- **Reasoning Transparency:** AKD systems must share workflow reasoning steps openly, enabling peers to evaluate underlying assumptions and identify potential limitations. In the near future, this will become an important research artifact needed for scholarly review.

Human Oversight and Scientific Accountability

Human oversight is a necessity in AKD. Every critical decision point in the scientific workflow must be designed to allow human scientists to intervene: to direct or redirect agents, inspect reasoning steps, override automated outputs, and independently validate results. The ultimate responsibility for interpreting and disseminating scientific conclusions must remain with human researchers.

AKD systems must be transparent by design to support the rigor of scientific inquiry. Opaque “black-box” systems are insufficient unless their reasoning is traceable and explainable. In particular, AI-generated outputs should include citations to peer-reviewed literature or validated datasets and support on-demand fact checking (Marinescu et al., 2025). Tools like FactReasoner demonstrate the feasibility of assessing factual accuracy by decomposing responses into atomic claims and verifying them against trusted sources.

Moreover, scientists must be informed about the limitations and assumptions behind the AI tools they use. AKD systems must clearly disclose any model constraints, uncertainties, or potential biases—especially in high-impact or sensitive research contexts. These systems should also be designed to identify contradictory findings, highlight underrepresented perspectives, and surface data gaps that could affect the validity of conclusions.

All outputs should be accompanied by structured logs and documentation that allow users to trace back reasoning processes, audit intermediate steps, and pinpoint potential sources of error.

Reproducibility and Provenance

Reproducibility is a cornerstone of scientific rigor, and this element of gold standard science is even more critical in AI-augmented discovery workflows. Within AKD, reproducibility should be understood as the ability of independent researchers to replicate results using the same data,

code, and workflow configurations. Although achieving perfect reproducibility with LLMs and other stochastic systems can be difficult due to inherent randomness, it is not unattainable.

With careful logging, rigorous version control, and consistent parameter constraints, LLM-based outputs can be made more reproducible. Therefore, reproducibility should be treated as a design goal within all AKD workflows.

To support this, AKD systems must have some standardized characteristics and outputs:

- Automated logging of datasets, model architectures, preprocessing steps, hyperparameters, and execution environments;
- Timestamped records of agent workflows, decisions, tool invocations, and human interactions; and
- Saveable and shareable workflows that can be re-executed under identical conditions and audited for process integrity.
- Standardized metadata schemas for research artifacts and its interactions with other components. These can be built on existing robust schemas such as the NASA's Unified Metadata Model (UMM) or the Planetary Data System (PDS4).

Agentic AI poses additional challenges due to its stateful and dynamic nature. The complexity underpinning agentic processes requires provenance tracking that goes beyond inputs and outputs to include the decision-making logic and reasoning paths followed by agents. Therefore, AKD systems should support the following aspects of provenance identification:

- Workflow-level explainability, allowing researchers to understand logic and flow;
- Integrated versioning across data, models, and tools; and
- Easy access to intermediate steps and reasoning logs to support review, reproducibility, and debugging.

These safeguards should not be retrofitted. Instead, safety must be built into the design from the outset. Trustworthiness should be an intrinsic design value, not an afterthought.

Together, these scientific trustworthiness imperatives form the foundation of the AKD framework as a responsible, science-centric AI approach. By embedding these conditions early into the design and development of AKD systems, institutions like NASA can lead the way in building a future where accelerated discovery is also ethical, transparent, and reproducible, upholding the core principles of scientific inquiry in the age of AI.

Advancing Accelerated Knowledge Discovery at NASA

AKD is uniquely positioned to enable several of NASA's science priorities including advancing scientific discoveries and fostering a culture of innovation (NASA Science Plan). First, AKD will advance scientific knowledge by making it easier to explore complex, cross-disciplinary

questions such as the search for life elsewhere in the universe. This fundamental science question requires data, information and advances from each scientific division. Each division contributes as follows:

- Astrophysics: Understanding how planets form and how to find and study them around other stars
- Biological and Physical Science: biological and physical systems in extreme space environments to achieve scientific breakthroughs not possible on Earth.
- Earth Science: Understanding behind atmospheric emission measurements, which can be used to search for signs of life on other worlds.
- Heliophysics: addresses how stellar activity and stellar magnetospheres affect planetary atmospheres and climate.
- Planetary Sciences: understanding of how geologic processes on Mars and ocean worlds in our solar system might give rise to habitable environments (“NASA Science Plan”, n.d.)

AKD will make it easier for scientists to explore these complex, cross-disciplinary questions by providing integrated access to a vast data and information ecosystem. AKD will streamline work for domain-specific research as well. Adaptive, AI-enhanced knowledge discovery scientific workflows can adapt to new observations, evolving models, and new hypotheses, dynamically integrating them into ongoing research processes. Such a responsive, interconnected approach will not only accelerate the scientific process but also foster deeper interdisciplinary research across domains.

Second, AKD will foster innovation by systematically integrating NASA's rich data and information systems, multimodal foundation models, scalable cloud and high-performance computing (HPC) infrastructure, and robust knowledge representation frameworks. In addition, AKD will ensure that NASA's science data are accessible and usable to everyone. Through this innovative approach, NASA will move beyond using AI as merely an analytical tool and instead embrace an operational future where AI functions as integrated cognitive enhancement tool supporting human scientists. Finally, as an investment in a high intellectual risk/high impact project, AKD will provide a new approach towards scientific discovery that has the potential to transform scientific workflows.

Crucially, trust is the foundation upon which the AKD framework must be built. NASA holds a high degree of trustworthiness grounded in a long record of scientific excellence, transparency, collaboration and open principles. Given this relationship, AKD must incorporate the trustworthiness imperatives outlined previously as essential design principles. These include demonstrable system reliability, explainable AI reasoning, detailed provenance for data and workflows, proactive mechanisms for detecting and mitigating hallucinations and biases, and rigorous adherence to open science principles. The consistent use of verified, authoritative sources and the generation of transparent, auditable logs of all AI decision processes are indispensable for maintaining both scientific integrity and public confidence in AKD-driven research.

Ultimately, NASA’s proposed open, transparent, and rigorously reproducible agentic framework — designed from the ground up with trustworthiness imperatives at its core — will empower breakthrough discoveries across the full spectrum of NASA’s science portfolio. From deciphering the biology of the smallest cells to unraveling the mysteries of the largest galaxies, AKD will serve as an enabler of 21st century science.

Acknowledgements

This work was supported by the NASA Office of the Chief Science Data Officer. We acknowledge the AKD design and development team from NASA IMPACT and IBM Research for their contributions. Special thanks to Rachel Wyatt and Derek Koehl for proofreading, editing, and improving the manuscript.

Data Availability Statement

Data were not used nor created for this research.

Conflict of Interest Statement

The authors have no conflicts of interest to disclose.

References

[1] Ren, Shuo, Pu Jian, Zhenjiang Ren, Chunlin Leng, Can Xie, and Jiajun Zhang. “Towards Scientific Intelligence: A Survey of LLM-Based Scientific Agents.” arXiv, April 17, 2025.

<https://doi.org/10.48550/arXiv.2503.24047>.

[2] Dawid, Herbert, Philipp Harting, Hankui Wang, Zhongli Wang, and Jiachen Yi. “Agentic Workflows for Economic Research: Design and Implementation.” arXiv, April 13, 2025.

<https://doi.org/10.48550/arXiv.2504.09736>.

[3] Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune, and David Ha. "The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search." arXiv, April 10, 2025. <https://doi.org/10.48550/arXiv.2504.08066>.

[4] Moss, Adam. "The AI Cosmologist I: An Agentic System for Automated Data Analysis." arXiv, April 4, 2025. <https://doi.org/10.48550/arXiv.2504.03424>.

[5] Yannis, Ioannidis. "The 5th Paradigm: AI-Driven Scientific Discovery." *Communications of the ACM*, n.d. <https://cacm.acm.org/opinion/the-5th-paradigm-ai-driven-scientific-discovery/>.

[6] Gray, Jim. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Edited by Tony Hey, Stewart Tansley, and Kristin Tolle. Microsoft Research, 2007.

[7] NASA, "NASA's Science Vision- NASA Science." Accessed: Jul. 22, 2025. [Online]. Available: <https://science.nasa.gov/about-us/smd-vision/>

[8] NASA, "Your Gateway to NASA Earth Observation Data | NASA Earthdata." [Online]. Available: <https://www.earthdata.nasa.gov/>

[9] NASA, "Heliophysics Data Portal." [Online]. Available: <https://heliophysicsdata.gsfc.nasa.gov/websearch/dispatcher>

[10] NASA, "Planetary Data Ecosystem | Planetary Data Ecosystem." [Online]. Available: <https://planetary.data.nasa.gov/>

[11] NASA, "NASA Open Science Data Repository." [Online]. Available: <https://www.nasa.gov/osdr/>

[12] "IRSA," NASA/IPAC Infrared Science Archive. [Online]. Available: <https://irsa.ipac.caltech.edu/frontpage/>

[13] NASA, "Science Discovery Engine." [Online]. Available: <https://sciencediscoveryengine.nasa.gov/app/nasa-sba-smd/#/home>

[14] Marinescu, Radu, Debarun Bhattacharjya, Junkyu Lee, Tigran Tchrakian, Javier Carnerero Cano, Yufang Hou, Elizabeth Daly, and Alessandra Pascale. "FactReasoner: A Probabilistic Approach to Long-Form Factuality Assessment for Large Language Models." arXiv, February 25, 2025. <https://doi.org/10.48550/arXiv.2502.18573>.

[15] Ramachandran, R., Bugbee, K., & Murphy, K. (2021). From open data to open science. *Earth and Space Science*, 8, e2020EA001562. <https://doi.org/10.1029/2020EA001562>

[16] NASA Open Science Training Team. (2025, July 10). NASA Open Science Essentials. Zenodo. <https://doi.org/10.5281/zenodo.15857836>

[17] NASA Science Plan: A Vision for Scientific Excellence.
<https://assets.science.nasa.gov/content/dam/science/cds/about-us/2025/2025-2026-NASA-Science-Plan.pdf>