

# A2SD: Accelerating Scientific Innovation through Autonomous Discovery Systems

Michela Taufer<sup>1</sup>[0000-0002-0031-6377], Rafael Ferreira da Silva<sup>2</sup>[0000-0002-1720-0928], Benjamin Mintz<sup>2</sup>[0000-0002-4054-1229], Milad Abolhasani<sup>3</sup>[0000-0002-8863-3085], Rosa M. Badia<sup>4</sup>[0000-0003-2941-5499], Ewa Deelman<sup>5</sup>[0000-0001-5106-503X], Robert G. Moore<sup>2</sup>[0000-0002-1608-5411], and John Shalf<sup>6</sup>[0000-0002-0608-3690]

<sup>1</sup> University of Tennessee, Knoxville, TN, USA

<sup>2</sup> Oak Ridge National Laboratory Oak Ridge, TN, USA

<sup>3</sup> North Carolina State University, Raleigh, NC, USA

<sup>4</sup> Barcelona Supercomputing Center, Barcelona, Spain

<sup>5</sup> University of Southern California, Marina del Rey, CA, USA

<sup>6</sup> Lawrence Berkeley National Laboratory, Berkeley, CA, USA

**Abstract.** The 2025 Advancing Autonomous Scientific Discovery (A2SD) workshop convened researchers from academia, national laboratories, and industry to explore the transformative role of autonomy in scientific discovery. The workshop highlighted a convergence of artificial intelligence, robotics, and computational workflows into autonomous systems capable of accelerating the scientific process. Presentations and discussions spanned autonomous experimentation, intelligent workflow orchestration, digital twins, and agent-based systems for managing complex research ecosystems. Key challenges discussed included interoperability across heterogeneous infrastructures, near real-time data management under FAIR principles, reproducibility, and the integration of human oversight. The workshop also emphasized the need for modular software interfaces, federated learning models, and education initiatives to support a next-generation scientific workforce.

**Keywords:** Autonomous Scientific Discovery · AI-Driven Workflows · Cyberinfrastructure Interoperability.

---

This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

## 1 Introduction

The first Workshop on Advancing Autonomous Scientific Discovery (A2SD-2025) [1], held at the ISC High Performance Conference, marked a critical milestone in the evolution of scientific research. Against the backdrop of increasing complexity in modern scientific challenges and the growing volume of data from advanced instruments, the workshop brought together leaders from academia, national laboratories, and industry to explore how autonomous systems are reshaping the landscape of discovery. The aim was to assess the state of the art, identify common challenges, and foster cross-disciplinary dialogue toward the development of scalable, interoperable, and intelligent autonomous infrastructures.

Autonomous science [7] represents a paradigm shift that blends artificial intelligence (AI), robotics, and computational workflows into a cohesive framework capable of accelerating scientific cycles. This transformation promises to shorten the path from hypothesis to validation by enabling systems that can reason, adapt, and act with minimal human intervention. Participants at A2SD-2025 emphasized the value of tightly coupled loops between theory, experiment, and computation, where autonomous agents can drive discovery by optimizing experimental parameters, orchestrating workflows across facilities, and learning from previous results. The workshop demonstrated that, while progress has been substantial, realizing the full potential of autonomous discovery requires addressing persistent challenges in integration, reproducibility, and human-AI collaboration.

A2SD-2025 presentations highlighted a broad range of use cases, ranging from edge-to-cloud computational workflows to autonomous synthesis platforms and digital twin-enhanced experimentation. Discussions emphasized how experimental and computational domains, often seen as distinct, are increasingly converging. Common technical themes such as scheduling, fault tolerance, and resource elasticity now span both domains. Several discussions underscored the need for systems that can operate across heterogeneous infrastructures, from local instruments and edge devices to high-performance computing (HPC) and cloud environments. Portability, abstraction, and co-design emerged as foundational principles for building resilient and scalable autonomy.

In addition to technical depth, the workshop also emphasized broader ecosystem-level considerations. These include standardization of agent interfaces, mechanisms for federated learning and model sharing, and protocols for handling sensitive data and intellectual property. The discussions expanded into education, workforce development, and policy, recognizing that building an autonomous science infrastructure is not only a technical endeavor but also a cultural and organizational transformation. The need to democratize access to advanced experimentation and foster fair participation across institutions was framed as essential to the future of autonomous discovery.

This paper synthesizes the key insights, findings, and community-driven recommendations from A2SD-2025. It draws upon the workshop presentations, panel discussions, and shared notes to provide a brief state-of-the-art review.

## 2 Technological Foundations and Advances in Autonomous Scientific Discovery

Several presentations showcased how autonomous experimentation is rapidly transforming the pace and scale of scientific discovery. Speakers emphasized the transition from domain-specific automation to modular, experiment-agnostic platforms. These systems incorporate digital twins, virtual representations of physical experiments, to enable near real-time monitoring, simulation, and optimization. Prof. Milad Abolhasani highlighted how digital twins are used not only to guide decision-making in microfluidics and chemical synthesis but also to reduce the reliance on trial and error by simulating outcomes before physical execution [6]. Reinforcement learning was presented as a key technique for rapidly identifying optimal experimental policies, with AI agents generating their training data to fine-tune decision strategies. The integration of miniaturized systems and robotic platforms further accelerates iteration cycles and reduces material waste.

In parallel with experimental advances, the workshop featured multiple contributions focused on the automation of computational workflows. Dr. Rosa M. Badia outlined how workflow engines, such as PyCOMPSs [2], enable autonomous execution in heterogeneous environments, including HPC clusters, cloud services, and edge devices [4]. These systems support runtime decisions for scheduling, failure handling, and resource reallocation, key capabilities for scalable and resilient scientific computing. Features such as fault tolerance policies, dynamic resource elasticity, task checkpointing, and hardware-aware task constraints enable workflows to adapt to changing execution contexts without requiring human intervention. The ability to orchestrate computations across an edge-to-cloud continuum was seen to be highly synergistic with experimental autonomy, creating opportunities for seamless integration of data acquisition, processing, and interpretation.

Several talks advanced a vision of AI agents that go beyond automation to actively contribute to scientific reasoning. Dr. Rob Moore and Dr. John Shalf emphasized the role of agentic AI in managing complex and distributed research ecosystems. These agents act as digital scientific assistants, helping to automate hypothesis testing, monitor experimental states, and recommend next steps based on historical data and learned models. They operate in physical labs, edge devices, and supercomputers, acting as intelligent intermediaries that can coordinate workflows and adapt strategies based on near real-time data. Real-world applications presented at the workshop illustrated how intelligent agents are being used to coordinate multi-stage scientific processes, manage resources across facilities, and support adaptive experimentation in various domains. This agent-based paradigm requires deep co-design that spans software, hardware, data, and experimental protocols to enable interoperability, reuse, and reproducibility across institutions.

An emerging theme across the presentations was the tight coupling of experimental and computational workflows to form autonomous discovery loops. Prof. Ewa Deelman and Prof. Michela Taufer presented compelling examples

where high-level workflow abstractions are compiled into executable pipelines that span multiple facilities, ensuring provenance and reproducibility throughout the process. These workflows must navigate challenges such as asynchronous data flows, heterogeneous instrumentation, and policy-based access constraints. Abstraction layers, such as those implemented in Pegasus [3] and PyCOMPSs, have been shown to facilitate the integration of diverse components, allowing scientists to focus on domain-specific questions. National initiatives, such as the National Science Data Fabric (NSDF) [8][5], aim to provide a shared cyberinfrastructure that connects distributed instruments, storage, and compute environments. NSDF is an example of a software ecosystem that enables live autonomous steering of a neutron diffraction experiment by integrating near real-time data streaming, persistent storage, and interactive dashboards with Bayesian active learning methods to dynamically adjust measurement parameters and optimize scientific discovery. These capabilities align with the goals of autonomous science by enabling near real-time data curation, metadata enrichment, and federated access to AI-ready datasets across institutions.

A shared emphasis on reproducibility in both computational and experimental domains highlighted the need for standardized metadata, robust fault handling strategies, and transparent decision making to ensure the credibility and transferability of results from autonomous science. The workshop also raised concerns about reproducibility challenges across geographic and institutional boundaries, where differences in data formats, experimental setups, and regulatory constraints can hinder consistent replication of results.

### 3 Cross-Cutting Challenges and Collaborative Opportunities in Autonomous Science

During the discussion sessions, the participants explored the tension between the desire for standardization and the variety of experimental and computational environments. Although universal standards could facilitate broader integration, institutional constraints and unique hardware-software combinations often require flexible and layered approaches. Several speakers noted that full standardization may be unrealistic, but shared abstractions and modular interfaces could offer a practical path forward. Abstraction layers were proposed as a way to enable interoperability across heterogeneous systems without imposing rigid constraints. Participants also reflected on past efforts to generalize software across domains, recognizing that although full generalization was often not achievable, many underlying patterns and design principles could still be reused and adapted to new settings. The discussions acknowledged that many laboratories rely on custom-built instruments with proprietary interfaces, underscoring the need for adaptable middleware that can abstract hardware-specific details while facilitating workflow interoperability.

Data management has emerged as one of the most pressing and complex issues facing autonomous science. The discussions emphasized the importance of applying the FAIR principles [9] in near real-time, from data collection to

long-term storage and reuse. The participants discussed the value of federated data architectures, which allow each institution to maintain control while enabling cross-facility collaboration. There was strong support for tools that can automatically capture metadata and assess data quality, especially given the scale and velocity of modern experimental systems. The topic of model sharing also featured prominently, with federated learning identified as a promising technique to move models instead of data. This approach not only enhances privacy and security but also supports scalable collaboration and decentralized innovation. Participants also emphasized the importance of publishing negative results and retaining rich metadata from autonomous workflows, recognizing that failed experiments can provide valuable training data and help prevent redundant exploration.

The role of human oversight in autonomous workflows was a recurring theme throughout the discussions. While autonomy can reduce manual workload and improve efficiency, participants stressed the need for well-defined protocols to determine when and where human input should remain essential. Examples from both experimental and computational workflows illustrated the risks of removing human judgment too early, especially when safety, credibility, or ethics are involved. Ethical considerations included the risk of indoctrinated science driven solely by data, without adequate theoretical grounding, and the potential propagation of bias in AI models trained on unbalanced or simulated datasets. There was also concern about overreliance on AI systems without a clear understanding of their limitations. To address these concerns, the group emphasized the need for transparent decision-making processes, validation mechanisms, and frameworks that allow scientists to inspect and override automated actions when necessary.

Participants agreed that technological progress must be accompanied by thoughtful investment in education and community development. Training the next generation of scientists to work alongside autonomous systems will require rethinking traditional curricula. Suggestions included integrating AI and data literacy into scientific education, developing hands-on training with virtual or simulated laboratories, and creating interdisciplinary programs that blend physical sciences, computing, and ethics. Beyond education, there was a strong call to democratize access to autonomous laboratory capabilities. This involves creating infrastructure that supports collaboration across institutions with varying resources and ensuring that scientific contributions across communities are recognized and valued. Community-driven initiatives and open platforms were identified as crucial tools for promoting equitable participation and advancing shared progress.

**KEY RECOMMENDATIONS:**

- Develop modular and interoperable software interfaces that support integration across heterogeneous experimental and computational systems.
- Implement federated data management and model sharing frameworks that enforce the FAIR principles and protect data privacy.
- Establish clear human-in-the-loop protocols to ensure scientific oversight and accountability in autonomous workflows.
- Redesign scientific education to include AI, data management, and ethics as core components of training for future researchers.
- Promote an open and collaborative infrastructure to democratize access to autonomous research tools and broaden community participation.

## 4 Conclusion

The A2SD-2025 workshop highlighted both the promise and complexity of building autonomous scientific discovery systems that can operate in various experimental and computational environments. Through a series of presentations and in-depth discussions, participants identified critical opportunities to advance autonomy in science, including the integration of intelligent agents, the coupling of digital twins with near real-time decision-making, and the development of resilient and reproducible workflows. Equally important were the cross-cutting challenges related to data management, system interoperability, human oversight, and fair access. As this field continues to evolve, long-term collaboration will require collaboration between disciplines, institutions, and communities. The insights and recommendations of this workshop offer a foundation for shaping a future in which autonomous systems accelerate discovery while preserving scientific rigor, open access, and trust.

## Acknowledgments

We gratefully acknowledge all A2SD-2025 workshop participants for their valuable presentations, insights, and discussions. This research used resources of the Oak Ridge Leadership Computing Facility, and is sponsored by the INTERSECT Initiative as part of the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725 and Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The material presented in the workshop is based on work supported by the National Science Foundation under Grant No. #2449103, #2513101, #2331152, #2223704, #2138811, #2103845. BSC authors acknowledge the partial support of projects PID2019-107255GB, CEX2021-001148-S, and PID2023-147979NB-C21 from the MCIN/AEI and MI- CIU/AEI /10.13039/501100011033 and by FEDER, UE, and by the Departament de Recerca i Universitats de la Generalitat de Catalunya, research group MPiEDist (2021 SGR 00412).

## References

- 1st Advancing Autonomous Scientific Discovery (A2SD-2025) Workshop. <https://autonomousscience.org/workshops/a2sd-2025/> (2025)
2. Badia, R.M., Conejero, J., Ejarque, J., Lezzi, D., Lordan, F.: Pycompss as an instrument for translational computer science. *Computing in Science & Engineering* **24**(2), 79–84 (2022)
3. Deelman, E., Vahi, K., Juve, G., Rynge, M., Callaghan, S., Maechling, P.J., Mayani, R., Chen, W., Ferreira da Silva, R., Livny, M., et al.: Pegasus, a workflow management system for science automation. *Future Generation Computer Systems* **46**, 17–35 (2015)
4. Lordan, F., Casas-Moreno, X., Cummins, P., Conejero, J., Badia, R.M., Sirvent, R.: Taming the Swarm: A Role-Based Approach for Autonomous Agents. In: European Conference on Parallel Processing. pp. 15–25. Springer (2025)
5. Luettgau, J., Martinez, H., Olaya, P., Scorzelli, G., Tarcea, G., Lofstead, J., Kirkpatrick, C., Pascucci, V., Taufer, M.: Nsdf-services: Integrating networking, storage, and computing services into a testbed for democratization of data delivery. In: Proceedings of the IEEE/ACM 16th International Conference on Utility and Cloud Computing. pp. 1–10 (2023)
6. Sadeghi, S., Mattsson, K., Glasheen, J., Lee, V., Stark, C., Jha, P., Mukhin, N., Li, J., Ghorai, A., Orouji, N., et al.: A self-driving fluidic lab for data-driven synthesis of lead-free perovskite nanocrystals. *Digital Discovery* (2025)
7. Ferreira da Silva, R., Abolhasani, M., Antonopoulos, D.A., Biven, L., Coffee, R., Foster, I.T., Hamilton, L., Jha, S., Mayer, T., Mintz, B., et al.: A grassroots network and community roadmap for interconnected autonomous science laboratories for accelerated discovery. *arXiv preprint arXiv:2506.17510* (2025)
8. Taufer, M., Martinez\*, H., Luettgau\*, J., Whitnah, L., Scorzelli, G., Newel, P., Panta, A., Bremer, T., Fils, D., Kirkpatrick, C.R., Pascucci, V.: Enhancing Scientific Research with FAIR Digital Objects in the National Science Data Fabric. *IEEE Computing in Science and Engineering (CiSE)* **25**(5), 39–47 (2023), 10.1109/MCSE.2024.3363828
9. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da Silva Santos, L.B., Bourne, P.E., et al.: The fair guiding principles for scientific data management and stewardship. *Scientific data* **3**(1), 1–9 (2016)