

## **DISCLAIMER**

**This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.**

# Toward Next-Generation Ecosystems for Scientific Computing

Harnessing community, software, and AI for cross-disciplinary team science

September 2025



```
xflux(F01DENS_FLUX:F05ENER_FLUX,i,j,k) = &
xflux(F01DENS_FLUX:F05ENER_FLUX,i,j,k) &
+cvisc*(/U(DENS_VAR,i-1,j,k) -U(DENS_VAR,i,
,U(DENS_VAR,i-1,j,k)*U(VELX_VAR,i-1,j,k)-U(DENS_VAR,i,
,U(DENS_VAR,i-1,j,k)*U(VELY_VAR,i-1,j,k)-U(DENS_VAR,i,
,U(DENS_VAR,i-1,j,k)*U(VELZ_VAR,i-1,j,k)-U(DENS_VAR,i,
,U(DENS_VAR,i-1,j,k)*U(ENER_VAR,i-1,j,k)-U(DENS_VAR,i,
_USM_MHD) || defined(FLASH_UGLM_MHD)
xflux(F07MAGY_FLUX:F08MAGZ_FLUX,i,j,k) = &
xflux(F07MAGY_FLUX:F08MAGZ_FLUX,i,j,k) &
+cvisc*(/U(MAGY_VAR,i-1,j,k) -U(MAGY_VAR,i,
,U(MAGZ_VAR,i-1,j,k) -U(MAGZ_VAR,i,
```



**Cover image:** The cover image represents the deflagration phase in a Type-Ia supernova in a Flash/Flash-X simulation. Image credit: Brad Gallagher.

## **About Argonne National Laboratory**

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Lemont, Illinois 60439. For information about Argonne and its pioneering and technology programs, see [www.anl.gov](http://www.anl.gov).

## **DOCUMENT AVAILABILITY**

**Online Access:** U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (<http://www.osti.gov/>), a service of the US Dept. of Energy's Office of Scientific and Technical Information.

**Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):**

U.S. Department of Commerce  
National Technical Information Service  
5301 Shawnee Road  
Alexandria, VA 22312  
**[www.ntis.gov](http://www.ntis.gov)**  
Phone: 800-553-NTIS (6847) or 703-605-6000  
Fax: 703-605-6900  
Email: [orders@ntis.gov](mailto:orders@ntis.gov)

**Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):**

U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831-0062  
**[www.osti.gov](http://www.osti.gov)**  
Phone: 865-576-8401  
Fax: 865-576-5728  
Email: [reports@osti.gov](mailto:reports@osti.gov)

## **Disclaimer**

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

**Report of the 2025 Workshop on  
Next-Generation Ecosystems for Scientific Computing:  
Harnessing Community, Software, and AI for Cross-Disciplinary Team Science**

**September 26, 2025**

**Workshop Location:**  
Chicago, IL

**Workshop Dates:**  
April 29 – May 1, 2025

**Suggested Citation:** Report of the 2025 Workshop on Next-Generation Ecosystems for Scientific Computing: Harnessing Community, Software, and AI for Cross-Disciplinary Team Science. L.C. McInnes, D. Arnold, P. Balaprakash, M. Bernhardt, B. Cerny, A. Dubey, R. Giles, D.W. Hood, M.A. Leung, V. López-Marrero, P. Messina, O.B. Newton, C. Oehmen, S.M. Wild, J. Willenbring, L. Woodley, T. Baylis, D.E. Bernholdt, C. Camaño, J. Cohoon, C. Ferenbaugh, S.M. Fiore, S. Gesing, D. Gómez-Zará, J. Howison, T. Islam, D. Kepczynski, C. Lively, H. Menon, B. Messer, M. Ngom, U. Paliath, M.E. Papka, I. Qualters, E.M. Raybourn, K. Riley, P. Rodriguez, D. Rouson, M. Schwalbe, S.K. Seal, Ö. Sürer, V. Taylor, and L. Wu. Report ANL-25/47, 2025.



## **Report Authors:**

Lois Curfman McInnes, Argonne National Laboratory  
Dorian Arnold, Emory University  
Prasanna Balaprakash, Oak Ridge National Laboratory  
Mike Bernhardt, Team Libra  
Beth Cerny, Argonne National Laboratory  
Anshu Dubey, Argonne National Laboratory  
Roscoe Giles, Boston University  
Denice Ward Hood, University of Illinois Urbana-Champaign  
Mary Ann Leung, Sustainable Horizons Institute  
Vanessa López-Marrero, Stony Brook University  
Paul Messina, Argonne National Laboratory (retired)  
Olivia B. Newton, University of Montana  
Chris Oehmen, Pacific Northwest National Laboratory  
Stefan M. Wild, Lawrence Berkeley National Laboratory  
Jim Willenbring, Sandia National Laboratories  
Lou Woodley, Center for Scientific Collaboration and Community Engagement  
Tony Baylis, Lawrence Livermore National Laboratory  
David E. Bernholdt, Oak Ridge National Laboratory  
Chris Camaño, California Institute of Technology  
Johannah Cohoon, Lawrence Berkeley National Laboratory  
Charles Ferenbaugh, Los Alamos National Laboratory  
Stephen M. Fiore, University of Central Florida  
Sandra Gesing, US Research Software Engineers Association  
Diego Gómez-Zará, University of Notre Dame  
James Howison, University of Texas at Austin  
Tanzima Islam, Texas State University  
David Kepczynski, Ford Motor Company  
Charles Lively, Lawrence Berkeley National Laboratory  
Harshitha Menon, Lawrence Livermore National Laboratory  
Bronson Messer, Oak Ridge National Laboratory  
Marieme Ngom, Argonne National Laboratory  
Umesh Paliath, GE Aerospace Research  
Michael E. Papka, Argonne National Laboratory, University of Illinois Chicago  
Irene Qualters, Los Alamos National Laboratory (retired)  
Elaine M. Raybourn, Sandia National Laboratories  
Katherine Riley, Argonne National Laboratory  
Paulina Rodriguez, The George Washington University  
Damian Rouson, Lawrence Berkeley National Laboratory  
Michelle Schwalbe, National Academies of Sciences  
Sudip K. Seal, Oak Ridge National Laboratory  
Özge Sürer, Miami University  
Valerie Taylor, Argonne National Laboratory  
Lingfei Wu, University of Pittsburgh

# Contents

|   |           |
|---|-----------|
| <b>Executive Summary: Reimagining Scientific Computing in the Age of AI</b>         | <b>ii</b> |
| <b>1 Significant Challenges Facing Computational Science</b>                        | <b>1</b>  |
| 1.1 Workshop objectives . . . . .   | 2         |
| 1.2 Potential impact . . . . .  | 3         |
| 1.3 Socio-technical co-design for next-generation scientific computing . . . . .    | 4         |
| <b>2 Breaking Down the Challenges</b>   | <b>6</b>  |
| 2.1 Software ecosystems for AI in scientific computing . . . . .                    | 6         |
| 2.2 Cross-disciplinary collaboration and AI for scientific software teams . . . . . | 8         |
| 2.3 Pedagogy and workforce development in the age of AI . . . . .                   | 10        |
| 2.4 Crosscuts and interdependencies . . . . .                                       | 11        |
| <b>3 Required Research Directions and Community Actions</b>                         | <b>12</b> |
| 3.1 Software ecosystems for AI in scientific computing . . . . .                    | 13        |
| 3.2 Cross-disciplinary collaboration and AI for scientific software teams . . . . . | 14        |
| 3.3 Pedagogy and workforce development in the age of AI . . . . .                   | 15        |
| 3.4 Required community actions . . . . .  | 16        |
| <b>4 Conclusion and Next Steps</b>  | <b>17</b> |
| <b>Acknowledgments</b>  | <b>18</b> |
| <b>References</b>   | <b>19</b> |
| <b>A Glossary</b>   | <b>25</b> |
| <b>B Workshop Description</b>   | <b>26</b> |
| <b>C Workshop Participants</b>  | <b>28</b> |
| <b>D Workshop Agenda</b>  | <b>30</b> |

## Executive Summary: Reimagining Scientific Computing in the Age of AI

The 2025 *Workshop on Next-Generation Ecosystems for Scientific Computing: Harnessing Community, Software, and AI for Cross-Disciplinary Team Science* convened in Chicago, IL, from April 29 to May 1, 2025. The event brought together more than 40 experts from high-performance computing (HPC), AI, computational science, software engineering, applied mathematics, social sciences, and community development to chart a path toward more powerful, sustainable, and collaborative scientific software ecosystems.

Workshop discussions underscored a powerful truth: scientific computing is at a turning point. As AI grows more capable and scientific questions become more complex, a transformation is unavoidable. The traditional ways we compute, code, and collaborate are no longer adequate. Periods of profound transformation bring both unprecedented opportunity and inherent uncertainty. Hence, addressing the complexity of this moment requires more than incremental progress—it necessitates strategic risk-taking and foundational shifts in approach. Just as the rise of digital computing reshaped nearly every scientific and industrial domain, AI-integrated software ecosystems must undergo that same level of conceptual reinvention, driven by rigorous inquiry and a willingness to redefine established frameworks.

We face the urgent need for a bold, enduring transformation—one that extends beyond technology alone. We envision a new kind of agile and robust scientific computing ecosystem built through *socio-technical co-design*—i.e, the intentional, integrated development of social and technical components as interdependent parts of a unified strategy. This approach integrates cutting-edge advances in AI and software—such as large language models (LLMs), scientific machine learning (SciML), reasoning models, and coding agents—with strategies to foster cross-disciplinary team building, training, and collaboration. Two key themes emerged:

- AI is not just another tool—it can dramatically boost scientific productivity, catalyzing new discovery. Moreover, while demonstrably transformative, AI's evolution is not known.
- Supporting this transformation in computational science will require a broad, cross-disciplinary community and sustained investment in preparing the next generation and empowering today's experts.

With these guiding themes, the workshop focused on three major challenges in future scientific computing:

- Advancing scientific software ecosystems for HPC and AI of today and tomorrow, while ensuring results remain scientifically valid.
- Fostering collaboration among scientists, AI experts, research software engineers, educators, community leaders, and even AI systems themselves—bridging traditionally siloed roles and domains. This includes private-sector partners, key to translating scientific research into real-world applications.
- Rethinking training and workforce development, to help today's and tomorrow's researchers adapt to rapid changes in tools and technology.

The workshop identified research directions and community actions to address these challenges. These include building modular, trustworthy, AI-powered scientific software systems; creating frameworks where humans and AI systems can develop and validate code together; overhauling training pipelines to reflect the pace and nature of modern computational science; promoting responsible innovation with guidelines for AI use; and launching pilot programs and partnerships that test new ways of learning and working.

In short, the workshop laid out a vision for the future: one where AI, software, hardware, and human expertise work hand in hand to accelerate scientific progress, broaden access, build the workforce, and preserve the integrity of the scientific method. This report marks the beginning of a long-term, community-driven effort to turn that vision into reality.

The workshop gave us a fresh, encouraging perspective. We must recognize that—more than ever—the future of science depends on more than computational power; it depends on how we embed human insight within intelligent systems. When we align AI's capabilities with the creativity, intuition, and discernment of researchers, we open new pathways for discovery. Purposeful integration of technical and societal dimensions creates a living system—one that learns, adapts, and accelerates progress where it matters most.

# 1 Significant Challenges Facing Computational Science

Throughout history, the most profound technological advances have challenged society’s capacity to adapt and respond at scale. Artificial intelligence (AI), while still immature, is already poised to influence nearly every facet of human endeavor (see, e.g., [1]). Similarly, while its full impact remains uncertain, AI’s disruptive and beneficial potential in science and engineering is clear. In turn, how we integrate AI across science and engineering may shape future global leadership.

Scientific discovery through computing underpins U.S. innovation, economic growth, and national security. Over the past fifty years, advances in computing have revolutionized how scientists learn, experiment, and theorize. High-performance computing (HPC)—including modeling and simulation, data analytics, machine learning, and AI—is a necessary means of discovery and innovation in essentially all areas of science, engineering, technology, and society [2, 3, 4, 5, 6, 7]. Today, that pursuit faces formidable challenges. As scientific data, modeling, simulation, and AI grow in volume, complexity, and strategic value, they have become central to the competitiveness of the nation’s industrial and manufacturing sectors such as automotive, aerospace, materials, and others. To remain a global leader, the U.S. must act with urgency and intention to elevate its scientific research infrastructure—especially through scientific software ecosystems that integrate AI and advanced computing capabilities. This is no longer a question of optional investment; it is a foundational requirement for ensuring our socioeconomic prosperity, maintaining global competitiveness, strengthening national security, and securing a higher quality of life for future generations.

Meeting this challenge requires broad national commitment and coordination. The scientific computing community cannot transform alone. Government, academia, national labs, and the private sector must invest together in the technical, institutional, and human systems that advance science. Scientific computing must evolve not only in tools and platforms but also in collaboration, governance, and talent development.

To meet this challenge, we must develop and adopt innovative approaches to the integration and application of scientific computing, learning from the past and preparing for the future. Scientific progress now hinges on our ability to integrate AI not only into our research methodologies but also into the underlying software ecosystems that power them [8, 9, 10, 11, 12, 13]. This integration must be safe, reliable, and agile, demanding a new generation of scientific software tools that are robust, scalable, and adaptable to rapidly evolving computing platforms—not retrofits, but purpose-built ecosystems capable of accelerating discovery in areas such as materials design, energy, biology, astrophysics, manufacturing, and security.

Equally important is the human dimension. Scientific advancement depends not only on tools and technology but also on the people who create and use them. We must cultivate a workforce of researchers and technologists who are proficient in developing and using emerging software ecosystems and AI-enabled research methodologies [7, 14, 15, 16, 17]. These people must be empowered to collaborate across disciplines, think computationally, and adapt rapidly to evolving technologies. Building this talent pipeline is not merely about education—it is about building the intellectual foundation for the nation’s scientific enterprise.

This is a multidimensional challenge—requiring urgent, coordinated action. We must advance technology, reimagine software, and cultivate a highly skilled workforce simultaneously. The United States aspires to remain a global leader in science and innovation [18], and progress in scientific computing is key to that vision. Over the coming decades, state-of-the-art science will increasingly rely on the effective, pervasive, and responsible use of AI, placing new demands on our tools, methods, and practices. A key enabler of this transformation will be the development of next-generation scientific software that is not only AI-integrated but also scalable, sustainable, and trustworthy. But the transformations needed for effective scientific progress are too large and pervasive to be done piecemeal by independent actors. The scope and urgency of the challenges demand coordinated, community-driven action across institutions, disciplines, and sectors—and ultimately, across national boundaries, while keeping the flexibility required for the dynamic evolution of scientific computing ecosystems. While these early efforts may focus on national coordination, lasting progress will require globally-connected community efforts.

## 1.1 Workshop objectives

To address these critical needs, the workshop *Next-Generation Ecosystems for Scientific Computing: Harnessing Community, Software, and AI for Cross-Disciplinary Team Science* was held in Chicago, IL, during April 29 - May 1, 2025.<sup>1</sup> The workshop convened a broad group of experts spanning HPC, AI, computational science, domain sciences, applied mathematics, software engineering, cognitive and social sciences, and community development. Participants collectively examined how we can co-design next-generation scientific software ecosystems to enable cutting-edge research, foster cross-disciplinary collaboration, and scale effectively across institutions and infrastructure. The workshop introduced the methodology of *socio-technical co-design for next-generation scientific computing*, explained in Section 1.3, as a complement to insights from other forward-looking community reports (e.g., [8, 9, 11, 19, 20, 21]). By intentionally blending technical and community factors, the workshop aimed to shape a future where thriving, cross-disciplinary communities drive the next wave of scientific discovery, with high-quality scientific software as a keystone of sustained collaboration and scientific progress.

**Scientific computing ecosystems.** In the context of this report, *ecosystems for scientific computing* are dynamic socio-technical systems made up of people, technologies, infrastructure, institutions, workflows, and cultural practices that collectively support the development and evolution of scientific computing. Ecosystems are role- and context-dependent—shaped by the interactions among stakeholders, researchers, developers, users, institutions, and AI systems. A successful ecosystem integrates technical components (such as software libraries, AI technologies, and computational platforms) with social dimensions (such as training, governance, and collaboration norms) to enable scientific progress. We believe that such ecosystems will emerge through the deliberate co-design of software platforms, scientific tools, community structures, and educational models—systems that support rapid change and are coordinated through well-designed, cooperative approaches rather than centralized control. Figure 1 illustrates aspects of these relationships.

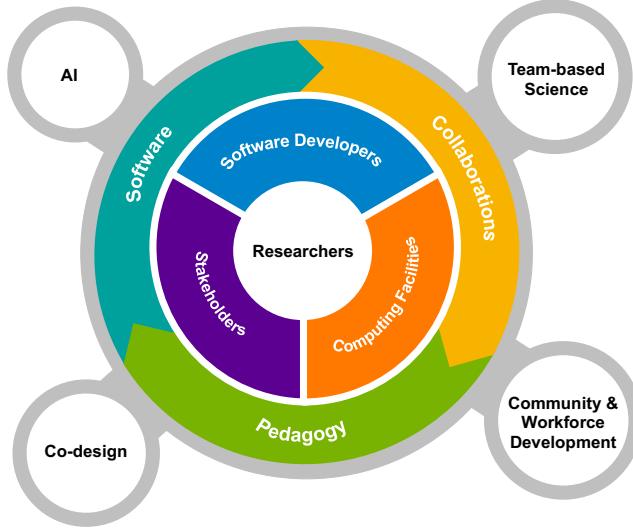


Figure 1: Toward next-generation ecosystems for scientific computing: As motivated by the needs of team-based science in an AI-driven future, we must advance software, cross-disciplinary collaboration, and pedagogy (outer ring of this figure) through co-design, all while developing workforce and community. This work demands strong collaboration among researchers (in science domains, applied math, computer science, HPC, AI, etc.), software developers, stakeholders, and computing facilities (inner rings of this figure).

<sup>1</sup> Appendices B, C, and D provide the workshop description, list of participants, and agenda, respectively.

## 1.2 Potential impact

Dynamic computational workflows, encompassing modeling, simulation, data analytics, and emerging AI approaches, as enabled by the instrumental building blocks of scientific software, are key to next-generation science [8, 22, 23, 24, 25, 26]. Ultimately, the promise—and eventual impact—of addressing the urgent challenges introduced above rests on the ability of distributed, cross-disciplinary teams to effectively integrate varied knowledge, perspectives, policies, requirements, and methods toward producing innovative solutions for scientific computing. Figure 2 illustrates the cross-disciplinary collaboration that is typical in advanced computational science—broadly considering simulations in physics, chemistry, materials science, biology, and so on,<sup>2</sup> which leverage heterogeneous computing architectures. These scientific codes build on low-level programming models and runtimes, libraries for math and visualization, tools to understand and optimize performance, emerging ML/AI technologies (such as distributed training frameworks, MLOps pipelines, and federated learning systems), application-specific components, and more.<sup>3</sup>

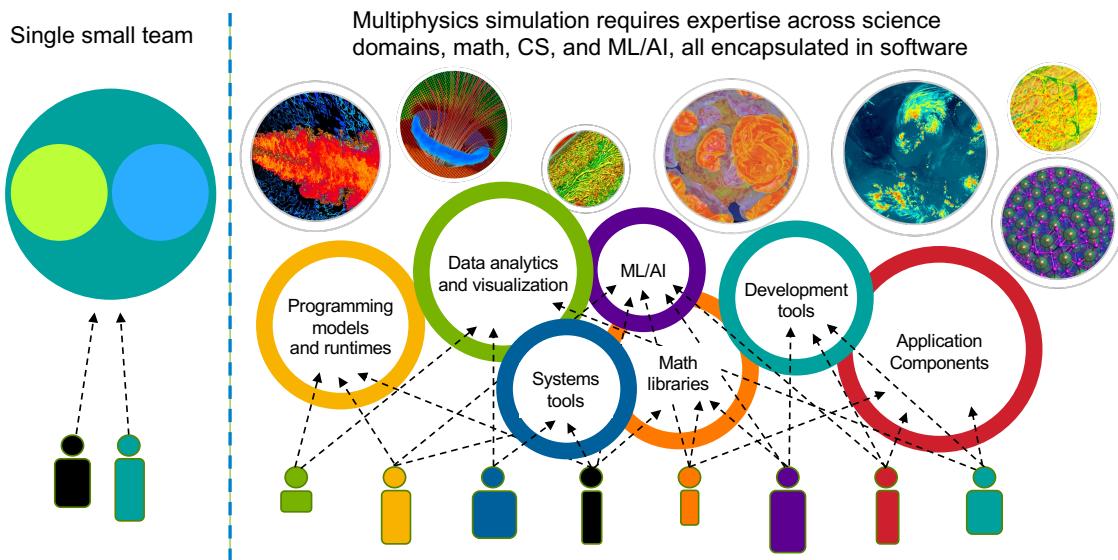


Figure 2: Simulations in advanced scientific computing (including materials science, astrophysics, nuclear energy, biology, engine design, weather prediction, and batteries, as represented by the science images of this figure) require collaboration across science domains, applied mathematics, computer science, ML/AI, and more, where high-quality software is a primary means of encapsulating expertise for use by others. Due to increasing science challenges and complexity, no longer can a single small team (shown on the left-hand side) independently develop all required functionality. Instead, reusable libraries and tools (shown on the right-hand side), developed by teams whose expertise spans across various topics, provide key functionalities that serve many applications. Dashed lines indicate multiple areas of work per person (e.g., the red-colored person in this diagram contributes to application components, development tools, and ML/AI capabilities).

Opportunities for advances via AI technologies abound, including code generation, automated design of experiments, AI agents, and real-time analysis during simulations. For example, a materials science team using AI to predict crystal structures needs seamless integration between calculations of density functional theory, ML training, and HPC scheduling, while building on programming models, systems tools, and so

<sup>2</sup>For example, see the wide range of scientific applications within the recent DOE Exascale Computing Project (ECP) [27, 28].

<sup>3</sup>For example, E4S [29] provides a foundational HPC-AI software ecosystem for science [30], including ECP libraries and tools [31, 32, 33] as well as popular AI packages.

on. Likewise, an astrophysics team using HPC and AI to study supernova explosions needs a multiphysics simulation code with shock hydrodynamics and nuclear burning, which might rely on math libraries for adaptive mesh refinement and on AI for parameterized physics models such as flame and equation of state.

Complexity arises due to heterogeneity across and within the various components constituting next-generation scientific computing ecosystems. This heterogeneity exists at many levels, for example in (1) hardware (e.g., CPUs, AI accelerators, quantum); (2) software (e.g., multitude of programming languages, environments, functionalities); (3) algorithms (e.g., mathematical, computer science, statistical, AI, domain science specific); (4) precision of operations (e.g., mixed vs. double precision, quantization in AI models); (5) data (e.g., storage, movement, processing, analysis, visualization); (6) computing environments (e.g., leadership computing facilities, federated, cloud, edge, quantum); and (7) the workforce (e.g., researchers, developers and users of software, domain scientists, management, stakeholders) or, more generally, cross-disciplinary teams, which in next-generation scientific computing may eventually include AI agents alongside humans [34, 35, 36]. This report considers the integration of AI into scientific software development as a transformative opportunity to accelerate discovery while addressing long-standing challenges in code quality, sustainability, and accessibility; other documents consider complementary topics (see, e.g., [37, 38, 39]). AI's promise lies not just in automating routine coding tasks but also in fundamentally reimagining how scientific software ecosystems evolve.

Collaboration is a primary mode of work in scientific computing, yet its potential is only beginning to be fully realized because of the complexity of cross-disciplinary research. The advent of AI, with its potential to influence human activities related to software development and scientific discovery, adds a transformative new dimension to this challenge. Next-generation ecosystems in scientific computing offer the promise of integrating broad domain expertise, methodological approaches, and AI-driven tools to bridge gaps between theoretical principles and practical implementations. These advances can enable more robust, maintainable, and scientifically accurate software, thereby accelerating scientific discovery.

We can build on experiences of innovative programs such as the U.S. Department of Energy's (DOE's) Exascale Computing Project (ECP; [27, 40]), Scientific Discovery through Advanced Computing (SciDAC) program [41], and Computational Science Graduate Fellowship (CSGF) program [42, 43]; the National Nuclear Security Administration's (NNSA's) Predictive Science Academic Alliance Program [44]; the National Science Foundation's (NSF's) Science and Technology Centers [45] and Engineering Research Centers [46]; Sustainable Research Pathways [47], and others. These initiatives have pioneered bold approaches for cross-disciplinary collaboration at scale and for cultivating both breadth across disciplines and depth within individual fields—countering traditional academic silos, as required for advances in scientific computing.

### 1.3 Socio-technical co-design for next-generation scientific computing

To tackle these urgent, complex, and intertwined challenges, the workshop introduced a novel methodology of **socio-technical co-design** for scientific computing, illustrated in Figure 3. The concept of co-design is widely used in many fields, referring generally to collaborative approaches to designing solutions by involving multiple stakeholders (especially those who will be using or are impacted by the outcomes) throughout the design process. For example, in the human-computer interaction community, co-design is strongly associated with participatory design. Co-design in HPC and AI has been critical to the design and implementation of contemporary computer architectures, applications, algorithms, and software, considering their interrelationships (as shown in the left-hand side of this figure) [48, 49].

Building on these co-design experiences, the workshop introduced the broader approach of socio-technical co-design, which embraces the intentional and integrated development of both technical components (e.g., software, AI, infrastructure) and social components (e.g., teams, institutions, practices, training) of scientific computing ecosystems. This expanded approach ensures that social and technical dimensions are addressed not in isolation but as tightly coupled elements of a unified, forward-looking strategy.

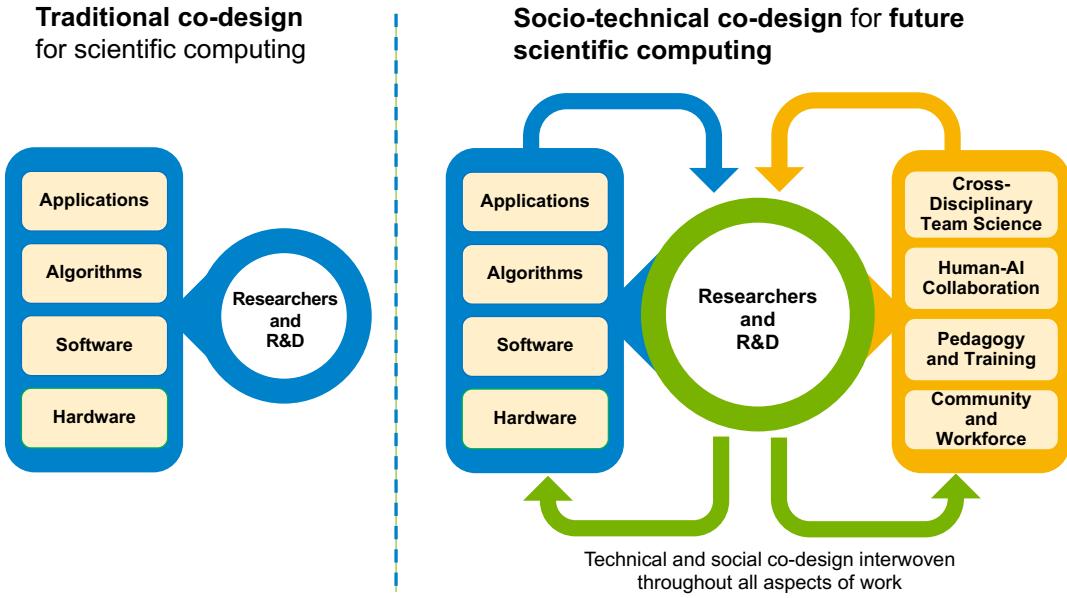


Figure 3: Socio-technical co-design for next-generation scientific computing intentionally interweaves technical and social elements throughout all aspects of work, while closely coupling cycles of R&D innovation between computing technologies and driving applications. By ensuring that technical and social dimensions are addressed as tightly coupled elements of a unified, forward-looking strategy, this holistic approach accelerates transformative impact across wide-ranging application domains.

**Report structure.** We organize insights from the workshop across three complementary axes that undergird scientific computing and urgently require better understanding and advancement to prepare for AI's growing role in shaping scientific discovery (see Figure 4):

- Software ecosystems for AI in scientific computing
- Cross-disciplinary collaboration and AI for scientific software teams
- Pedagogy and workforce development in the age of AI



Figure 4: Report structure.

Two crosscutting topics permeate all of these:

- AI as a catalyst to advance software productivity, collaboration, and pedagogy
- Community engagement and capacity building for next-generation scientific computing

These domains are deeply interconnected, requiring co-design of technical and social systems. Addressing them holistically presents an opportunity to foster robust and scalable research ecosystems capable of driving scientific discovery in the coming decade. We next delve into these challenges, including crosscuts and interdependencies, and we introduce priority research directions as well as actionable community strategies.

## 2 Breaking Down the Challenges

As scientific computing moves toward AI-driven, cross-disciplinary, community-integrated ecosystems, significant challenges arise in aligning technical tools with human needs and institutional capabilities. Sections 2.1 through 2.3 elaborate on challenges in the complementary topics introduced in Section 1, while Section 2.4 discusses crosscutting issues and interdependencies.

### 2.1 Software ecosystems for AI in scientific computing

We define a software ecosystem for AI-enhanced scientific computing as the entire software infrastructure that supports all computational steps involved in scientific discovery, ranging from system software at the lowest level to inference engines at the highest level. The inference engines rely upon data obtained from simulations, observations, archives, and other sources of scientific data. Simulation and analysis tools in turn depend upon middle layers of the software stack such as math and visualization libraries and workflow management tools. Tools to analyze, predict, and optimize performance, energy, and cost of simulations are also essential. The software ecosystems must not only integrate traditional numerical methods and scientific simulations but also accommodate rapidly evolving AI models, data pipelines, and heterogeneous computing environments (e.g., CPUs, GPUs, and quantum systems). While heterogeneity is not new to scientific computing, at present it is growing in complexity and ever-changing, raising questions about CPU-GPU memory transfers, mixed precision training, and quantum-classical hybrid algorithms. Key questions thus revolve around the urgency to identify what will be needed to design, develop, and maintain software that will successfully support next-generation science conducted under heterogeneous, dynamic, and fast-evolving environments, while simultaneously building the human enterprise and collaboration protocols. Specific software needs are discussed in the recent Report of the NSF/DOE Workshop on NAIRR Software [11], while software engineering challenges are presented in [20]. Creating extensible and sustainable ecosystems will require a strategic blend of technical innovation and community-driven co-design.

#### Key questions:

- How can we build extensible, traceable, and modular scientific software ecosystems that support both legacy and AI-generated components, and what degree of intrusion into existing code is needed and acceptable? Can AI assist in decomposing monolithic scientific codes into reusable components that support ecosystem integration? How can AI tools help with refactoring, modernization, and documentation of legacy scientific code?
- How can AI maintain scientific correctness while accelerating development? Scientific applications must adhere to physical laws, mathematical principles, and domain-specific constraints. How can we train AI systems to generate code that is scientifically valid and numerically stable?

- What validation frameworks are needed for AI-generated scientific code? Traditional software testing approaches may be insufficient for scientific applications where correctness extends beyond functional requirements to include adherence to scientific principles, uncertainty quantification, and reproducibility across different computational environments.
- How can standards for AI be designed and integrated into software, applications, tools, and reporting to ensure the technologies can be trusted and adopted by different segments of the scientific computing community with relative ease?

### Scientific challenges and opportunities:

**Fragmented ecosystem and interoperability challenges.** Scientific software ecosystems consist of wide-ranging codebases that already struggle with continually changing hardware paradigms and now must also integrate with modern AI frameworks. We face challenges in bridging traditional scientific computing environments with AI-driven development workflows, including deploying trained models in HPC environments and integrating scientific data formats with AI frameworks. Given that data is a critical component of AI-enabled scientific computing, and that computational results are often distributed across teams and siloed in storage systems, seamless access to data remains a major barrier. To support effective collaboration, we need standardized mechanisms for data sharing that address access control, provenance, and metadata. The lack of standardized interfaces, data formats, and integration protocols—especially for deeply integrated components like math libraries—further prevents seamless adoption of AI technologies and forces scientists to navigate multiple, incompatible systems. This fragmentation creates inefficiencies, requires manual intervention between simulation and analysis stages, and limits the ability to leverage AI across entire scientific workflows. Additional challenges relate to the need for performance optimization, such as model quantization, pruning, and knowledge distillation, as well as adaptability in dynamic AI landscapes.

**Reliability and validation gap in AI-generated scientific code.** A critical challenge is the disconnect between AI’s current capabilities and the requirements of scientific computing for reliability and replicability [50, 51]. Unlike general-purpose software, scientific applications must preserve physical principles, maintain numerical stability, and provide reproducible results across different computational environments. Current AI tools often generate syntactically correct code that may violate domain-specific constraints or introduce subtle numerical errors that compound over long simulations. This point is crucial because natural language specification is by its nature imprecise. Traditionally, programmers eliminate artifacts of imprecise specification through iterative verification and debugging. What can replace or accelerate this iterative convergence to trustworthy software in its development cycle remains an open question. This gap is exacerbated by the lack of standardized validation frameworks specifically designed for AI-generated scientific code, making it difficult to systematically assess whether generated solutions maintain both computational correctness [52] and scientific validity. Federated learning introduces additional challenges, including needs for privacy-preserving collaboration across institutions and addressing security threats such as adversarial attacks on scientific AI models and data poisoning. Additionally, human-in-the-loop and continuous learning mechanisms will remain essential for aligning AI-generated software with evolving scientific understanding.

**Development methodologies for hybrid modeling/simulation and AI.** While it will be impractical to develop entirely new scientific software ecosystems from scratch, reuse of existing software in these new modalities is likely to be challenging. The changes required to code will be deeply invasive; therefore, intrusion-aware design will become critical to minimize disruption when integrating new tools or AI systems into existing software. There is an ongoing debate in the scientific community about when and where traditional numerics should be used, where emerging scientific machine learning (SciML [53]) methods might replace them, and how the two can most effectively complement one other. We urgently need to

develop (and adopt) suitable metrics when evaluating SciML methods against traditional numerical methods [54, 55]. Middleware to assist in the creation and execution of automated computational workflows will continue to be essential but will now also need consideration for agentic AI systems [56]. To support scalable infrastructure growth, such advances require ecosystem-level funding models, broad coordination, and consortia involving national labs, academia, and the private sector.

## 2.2 Cross-disciplinary collaboration and AI for scientific software teams

Because the physical world is not divided into specific domains, those who seek to understand it must also eliminate unnatural boundary lines dividing their discourse. Computational science in general, and high-performance computing in particular, have pioneered advances in cross-disciplinary teams combining computing with science, engineering, and mathematical foundations. Such scientific teamwork, in which an interdependent set of group members share a common purpose, performance goals, and mutual accountability, has grown significantly to meet the needs of organizations and tackle formidable problems [57, 58, 59, 60]. However, the growing complexity of challenges in scientific computing demands increasing collaboration among researchers and potentially AI agents—where software is a primary means of collaboration.

Cross-disciplinary teams in scientific computing face communication and knowledge integration challenges due to diverse vocabularies, standards, and methodologies [60, 61]. These challenges are associated with the structure and norms of scientific research; the proliferation of data and computational technologies; the complexity of interactions; and the knowledge, skills, and attitudes present in teams. The rapid evolution of AI technologies complicates each of these issues and also introduces a unique set of challenges. A multipronged strategy addressing macro-, meso-, and micro-level factors (see Figure 5) that influence collaboration in scientific computing is of paramount importance; consequently, we structure the remainder of this section according to these factors.

### Macro-level factors (science and society): Structures and norms for collaboration

#### Key questions:

- How can existing funding structures be adapted to encourage greater and optimal interconnectedness in the scientific computing community? What new mechanisms, and possibly new sources for funding, are needed to encourage and maintain collaborations in scientific computing?
- How can existing norms for recognition and credit be adapted to encourage greater and optimal sharing and cross-disciplinary collaboration in the scientific computing community? What novel norms for recognition and credit in cross-disciplinary research are needed to overcome the limitations of siloed, unidisciplinary structures in academia?
- How can standards for AI technologies be designed to facilitate rather than constrain collaboration and innovation?

**Scientific challenges and opportunities:** The structures and norms that shape or otherwise enable scientific research are often at odds with the approaches and goals that characterize cross-disciplinary collaboration. These include the availability and distribution of funding (often insufficient for the full scope of

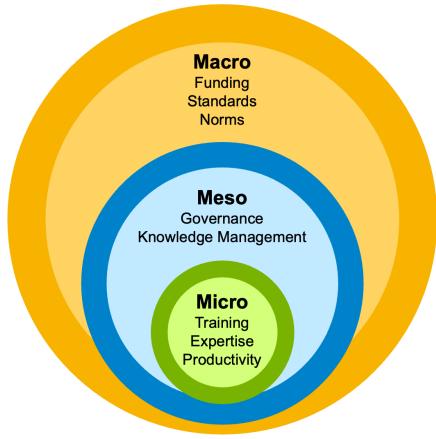


Figure 5: Team challenges span throughout macro-, meso-, and micro-level factors.

multi-institutional, cross-disciplinary multiteam systems [60] required for large-scale computing projects). Another macro-level challenge is norms for recognition and credit, which sometimes neglect or minimize the work of various types of contributors in projects, resulting in decreased motivation, satisfaction, and productivity. Reward systems need to adapt accordingly, reflecting a shared value system for all components of the evolving landscape, including the development and maintenance of the underlying supports such as software “plumbing,” project management, and communities that support collective needs.

### **Meso-level factors (organizations, interactions): Models of governance and knowledge management**

#### **Key questions:**

- What forms of governance are most beneficial for achieving fair and defensible decision making in cross-disciplinary collaborations in scientific computing?
- How can these governance models enable effective collaboration, accelerate innovation, and help ensure that all relevant perspectives are represented? What resources and tools are needed to enable the adoption of these governance models in multiteam collaborations?
- What approaches to knowledge management are beneficial for fostering trust, mitigating knowledge loss, and reducing barriers to contribution and collaboration?

**Scientific challenges and opportunities:** The complexity of collaboration in scientific computing is associated with the number of stakeholders and the differences in their requirements, approaches, and goals. Effectively managing this complexity is vital for the quality of collaboration processes and outcomes, including communication, coordination, decision-making, problem solving, performance, and knowledge transfer [60]. Access to knowledge must become ubiquitous, fair, and democratic, taking advantage of AI and cloud-based approaches. This challenge introduces an opportunity to envision models of governance and knowledge management that support distributed collaboration across institutions, ensuring shared responsibility for next-generation ecosystems in scientific computing.

### **Micro-level factors (teams and constituent members): Training for expertise and productivity**

#### **Key questions:**

- How can collaboration among various actors in scientific computing (including national laboratories, universities, and the private sector) be leveraged to design and provide opportunities for training in cross-disciplinary research models and teamwork competencies?
- How can human-centered strategies and AI technologies be designed to support the acquisition of *taskwork* competencies (e.g., knowledge, skills, and abilities required for writing high-quality software, evaluating and selecting software, modifying and adapting software) as well as *teamwork* competencies (e.g., knowledge, skills, and abilities required for effective communication, coordination, learning) in scientific computing?
- How can human-centered strategies and AI technologies be designed to minimize workload, reduce miscommunication, increase efficiency, and augment the capabilities of individuals and teams in scientific computing projects? How can human-centered strategies and AI technologies be designed to enhance team formation and team development across multiteam collaborations?

**Scientific challenges and opportunities:** Cross-disciplinary research teams face significant challenges in communication, which in turn can reduce their capacity for effectively utilizing collaborative member expertise, developing shared knowledge structures (i.e., shared mental models and transactive memory systems), and integrating their varied knowledge [61]. Because cross-disciplinary research demands the integration of

the varied, specialized expertise present in a team, its members must be adept in communication and exhibit openness toward one other—competencies and dispositions that influence the team’s expertise utilization and knowledge integration. Furthermore, individuals in these teams must develop expertise in topics deemed essential for scientific computing (e.g., reproducibility [51]). Both training and research are thus needed on attitudes, behaviors, and cognition that can enhance cross-disciplinary collaboration in the scientific computing community. Also needed are advances in AI-specific collaboration tools, such as shared Jupyter environments, collaborative model development platforms (GitHub for models), and automated code review systems for AI-generated code. For example, domain scientists need to be able to validate AI-generated finite difference schemes without understanding all details of the underlying transformer architecture.

**Multilevel factors: Trust in teammates, software, and AI.** Achieving an appropriate level of trust in teammates, software, and AI represents an additional significant challenge. Developing and maintaining trust across the macro, meso, and micro levels of collaborative scientific computing requires insightful leadership, time, repeated interaction, and intervention when trust is lost. A challenge is devising alternative means for achieving trust within an ecosystem—minimizing the workload required to develop trust and manage uncertainty associated with information (e.g., by leveraging wisdom of the crowds and collective intelligence). Also needed are trust protocols for AI agents, for example for AI agents to submit computational jobs to leadership-class facilities.

### 2.3 Pedagogy and workforce development in the age of AI

Pedagogy and workforce development in the computing sciences have always been challenging because of the need to understand and communicate ideas that span multiple disciplines. In academia an added challenge has been the difficulty of fitting in cross-disciplinary training in traditional department structures [5, 7]. Mission-driven research labs have been the natural homes for such efforts. With the advent of AI and growth of computational science—spanning beyond traditional roots in the physical sciences to include life sciences, social sciences, and virtually all aspects of science and society—there is more opportunity and urgency to integrate cross-disciplinary training. Many science and engineering departments will need to extend their curricula to include developing expertise in exploiting powerful computing platforms, including AI, in much the same way they have trained students in experimental methods in the past.

#### Key questions:

- How do we prepare the scientific workforce for an AI-augmented future? As AI tools become more sophisticated, the roles of domain scientists and software engineers will evolve. Critical questions include the following: What fundamental skills remain essential? What specific AI/ML concepts are essential? How do we maintain critical thinking and domain expertise when AI handles routine tasks? How do we ensure that the next generation can validate and improve upon AI-generated solutions? How can we ensure that AI models and datasets reflect a wide range of perspectives grounded in real-world contexts?
- How can strategies be developed to train and grow workforce communities when the requirements are unknown and rapidly changing? Moreover, how can we broadly enhance and evolve education to foster critical thinking, extend access, expand cross-disciplinary collaboration and communication, and create ubiquitous and democratic access to and implementation of AI for computational science? How can we effectively address community building and workforce development among various stakeholder groups?
- How can we build training catalogs for community-wide use, including all relevant perspectives, so that we can broaden the pool of scientists and engineers engaged in computational science?

### Scientific challenges and opportunities:

**Workforce development and knowledge transfer crisis.** Perhaps the most profound challenge identified in the workshop is the degree of uncertainty in how to prepare the future workforce for the disruptive changes underway in computing for science. As AI tools become more capable, there is growing concern about creating a generation of scientists who can use AI-generated code but lack the fundamental understanding needed to validate, debug, or extend it. This situation creates a paradox: AI tools designed to democratize software development may inadvertently create new barriers to deep scientific understanding. Simultaneously, the rapid pace of AI advancement makes it difficult for educational institutions and training programs to update curricula as technologies and practices change, while the shift toward AI-assisted development raises fundamental questions about what constitutes essential knowledge in the age of artificial intelligence. Building curricula and the workforce for rapidly changing technological ecosystems requires new instructional and workforce development paradigms, with higher reliance on apprenticeship models, including the social (group) learning that is supported by communities of practice [62].

**Strategies to bridge across disciplines.** The development of integrated, cross-disciplinary curricula in scientific computing has been a slow and uneven process; despite decades of progress, many academic institutions still lack programs in computational science and engineering. As next-generation scientific challenges increasingly demand complex, team-based approaches that span multiple disciplines, it is essential to accelerate efforts to design curricula that bridge scientific domains, computational methods, and technological modalities. To support this shift, there are growing opportunities to leverage cloud-based AI platforms and containerized environments that enable consistent, scalable training across diverse institutional contexts. We also need to incorporate critical professional competencies often missing from current curricula, such as project management, agile development practices, effective communication, and collaborative teamwork.

**Gaps in workforce readiness and training.** One of the most persistent challenges in computational science is the limited pool of qualified candidates for recruitment—a problem compounded by the scarcity of formal academic programs in the field. Although various approaches have been attempted, they have not fully addressed the underlying gap in training pathways. Apprenticeship models for hands-on learning have shown great promise as a complementary approach to formal instruction. These are typically well-supported by associated communities of practice, where practitioners can continue to share knowledge and tactics in an ongoing, informal manner that supports continuous learning [63]. Another promising approach is to shift from solely recruiting individuals with comprehensive subject-matter expertise to also recruiting based on potential, followed by targeted, domain-specific training. However, this strategy introduces a new set of challenges: developing high-quality training content, keeping it up to date with evolving technologies, and delivering it at scale. These tasks require sustained resource investment and strong coordination across the broader scientific and educational communities.

## 2.4 Crosscuts and interdependencies

To fully understand the challenges and opportunities identified in Sections 2.1 through 2.3, it is essential to consider two crosscutting themes that emerged consistently throughout the workshop: (1) the transformative role of AI in improving scientific software productivity, collaboration, and pedagogy; and (2) the central importance of community engagement and capacity building in enabling and sustaining these transformations in the computing sciences. These crosscutting themes intersect with all three core challenge areas and are foundational to building resilient next-generation scientific computing ecosystems.

**AI as a catalyst for scientific productivity and sustainability.** Artificial intelligence is not just a topic unto itself—it is a force multiplier across software ecosystems, collaboration frameworks, and pedagogical approaches. Section 2.1 discusses the need for AI-integrated software ecosystems that are modular, traceable,

and scientifically valid. AI can play a key role in this work by helping to advance software productivity and sustainability, such as enabling intelligent code generation, automating routine tasks, and supporting decision-making across heterogeneous computing environments. AI also presents new challenges in terms of verification, trust, and intrusion-aware integration.

AI can promote collaboration by acting as a bridge across disciplinary boundaries. As discussed in Section 2.2, AI can support shared understanding by translating domain-specific terminology, managing distributed workflows, and facilitating knowledge retention. AI systems can also act as teammates—assisting with code reviews, optimizing team formation, and recommending learning pathways. These roles require careful governance to ensure they support rather than hinder team cohesion and trust.

As explored in Section 2.3, AI can also transform pedagogy by offering new modalities of instruction, such as AI tutors, generative feedback, and adaptive curriculum planning. However, these benefits must be balanced against the risk of over-reliance, which could erode critical thinking and domain expertise. AI-enhanced education systems must therefore be grounded in robust, fair frameworks and designed to complement (not replace) human mentorship and community support.

**Community engagement and capacity building as a strategic imperative.** As AI transforms the technical foundations of scientific computing, the long-term success of that transformation will depend equally on the strength of the communities that adopt, adapt, and sustain it. Section 2.1 highlights that successful integration of AI into software ecosystems requires more than technical proficiency—it demands fluency in interdisciplinary thinking, sound judgment, and collaborative practices. Building this capacity across communities is essential for maintaining and evolving shared software infrastructure.

As discussed in Section 2.2, the human dimension of collaboration is central. Scientific progress increasingly depends on teams with complementary skills, diverse perspectives, and broad participation. Capacity-building strategies—such as training programs, fellowships, internships, apprenticeships, and institutional partnerships—should be designed to foster cross-disciplinary fluency, resilience, and leadership. Community norms, as well as opportunities for professional development, must evolve to recognize and reward all contributors, including those in roles that have historically been undervalued.

Section 2.3 underscores the need to rethink education and community building considering these dynamics. Curriculum innovation, community-led training models, and broad access to computational and human resources are critical. Well-supported communities create networks of practice that encourage and reward innovation and different perspectives, share lessons learned, and mentor the next generation of scientists, developers, and educators.

Together, these two crosscutting themes—AI to advance scientific productivity and sustainability, and community engagement with capacity building for next-generation scientific computing—represent both the engine and the scaffolding of future scientific computing. They are deeply interdependent: AI systems must be designed, deployed, and governed by engaged, well-supported communities; and those communities must be equipped with the tools, practices, and shared knowledge that AI can help accelerate. Advancing one without the other will limit impact. Investing in both will accelerate progress and ensure that next-generation ecosystems for scientific computing are not only more powerful, but also widely accessible, trusted, and sustainable.

### 3 Required Research Directions and Community Actions

Considering the challenges and opportunities outlined in Section 2, the workshop identified research directions and community actions needed for accelerating progress in next-generation scientific computing. As discussed in Sections 3.1 through 3.3, each priority research area combines practical needs with aspirational goals, aiming to foster a productive interplay between human insight and AI capabilities while maintaining

a strong foundation of scientific rigor and community engagement. The recommended research directions emphasize key design principles: modularity, interoperability, trustworthiness, and extensive participation. The intent is to support both near-term implementation and long-term sustainability of emerging software and collaboration models for scientific computing. Section 3.4 outlines strategic community actions that are essential for supporting the software ecosystems, collaborative teams, and educational innovations discussed throughout this report. Taken together, the priority research directions and community actions define a comprehensive agenda for building thriving, resilient, and forward-looking scientific computing ecosystems.

### 3.1 Software ecosystems for AI in scientific computing

Next-generation scientific computing demands a radical rethinking of how software ecosystems are designed, developed, and sustained. As AI becomes more integrated into scientific workflows, software infrastructure must evolve to support new modes of discovery, while maintaining the rigorous standards of scientific correctness, reproducibility, and performance. This section identifies research priorities for building AI-integrated scientific software ecosystems that are modular, trustworthy, and adaptable. The topics below reflect a blend of foundational challenges and forward-looking opportunities, ranging from infrastructure and interoperability, to AI-assisted code generation, to robust validation methods that incorporate community knowledge and formal guarantees. Taken together, these efforts can help create a robust, adaptive foundation for software ecosystems that meet the scale, complexity, and scientific integrity required for the AI-powered future of discovery.

**Software and infrastructure for next-generation science.** Existing software infrastructure, although rich in abstractions, libraries, and tools, is inadequate for the multipronged challenges of next-generation scientific applications. The growing demands of high-fidelity models and the fragmentation of system software across AI-driven hardware platforms have created major obstacles for scientific teams. Libraries and tools may be underutilized because of the complexity of manually navigating multiple packages, and most science teams cannot feasibly manage the complexity of hybrid HPC/AI environments without new modes of support. Similar software challenges were articulated in the National Artificial Intelligence Research Resource (NAIRR) Software Workshop [11].

A promising direction is to design software substrates that enable AI/ML and traditional simulation codes to interoperate scalably and adapt to evolving workflows, platforms, and algorithms. At present we have little understanding of what such substrates might look like, but we can begin with research to determine and minimize the intrusiveness required to connect legacy scientific software with new AI tools, coordinate multiple coexisting ecosystems (e.g., AI agents interfacing with scientific tools), and analyze how software components interact in real workflows. Of particular importance is exploring the use of AI agents in end-to-end scientific code development workflows, including interactions with math libraries, visualization tools, debuggers, job schedulers, and HPC performance analysis tools. Key areas of research include energy-efficient implementations of numerical and AI methods, tools for software usage analytics, AI-assisted debugging, and strategies for hybridizing traditional solvers with SciML methods. Also needed are metrics to evaluate emerging SciML approaches against traditional numerical methods.

**Development of science-aware AI code generation systems.** A key research priority is the development of AI systems designed and trained specifically for scientific computing contexts. This requires more than adapting general-purpose LLMs; it calls for domain-specific training on high-quality scientific codebases, the incorporation of physics-informed constraints, and the use of validation datasets that emphasize correctness as well as functionality. Research should focus on developing hybrid approaches that combine traditional symbolic reasoning with neural methods, enabling AI systems to respect physical laws, dimensional analysis, and numerical stability requirements. Promising techniques include physics-informed neural networks (PINNs), neural operators, and differentiable programming, which offer pathways for AI-generated

code to honor physical laws, conserve energy, and maintain numerical stability. Systems must be capable of constrained code generation (e.g., ensuring energy conservation in molecular dynamics) and must integrate domain-specific knowledge graphs that encode scientific principles. Formal verification methods tailored to scientific applications are also essential. Key components include training datasets curated for scientific applications with verified correctness and development of specialized architectures that can reason about mathematical relationships and physical constraints. This research direction should also investigate methods for uncertainty quantification in AI-generated scientific code, enabling users to understand confidence levels and potential error propagation. We also need advances in AI interpretability, such as understanding AI decisions in safety-critical scientific applications. To integrate existing code into new development, research is needed on AI tools that can generate, refactor, and modernize scientific code while ensuring correctness and maintainability. Work is needed to use AI to extract knowledge from existing codebases and generate useful, accurate documentation for end users and contributors. AI evaluations of code bases can also be used to inform teams and leaders of risks in a code base (fragility, limited testing, etc.) and to help teams make decisions about rewriting and refactoring. We must also explore systems where human developers interactively train AI models in context, for example, by correcting recommendations.

**Community-integrated validation and continuous learning frameworks.** Given the critical importance of correctness in scientific software, research is needed to develop validation frameworks that leverage community knowledge, expert input, and formal guarantees. These systems should support automated generation of test cases grounded in scientific principles, including testing against known analytical solutions where available. Techniques such as property-based testing and metamorphic testing are especially valuable for scientific domains, enabling validation of AI-generated code even in the absence of exact reference solutions. For example, fluid dynamics code produced by AI must satisfy divergence-free velocity fields and conserve mass properties that can be checked independently of a specific output. Research should also explore adaptive “living” validation systems that evolve with scientific understanding and incorporate domain expertise as part of the AI training feedback loop. Peer review mechanisms tailored to code generation could facilitate expert vetting of outputs and structured incorporation of corrections into model refinement. These systems may include code that is validated against observations and experiments that cannot be replicated. The use of formal methods for system correctness [64] and the verification of implementations of numerical methods [65], as well as AI methods [66, 67, 68], is needed. In addition, as AI-generated code and integrated workflows become more complex and widely adopted, research should address emerging security concerns, including software vulnerabilities and ecosystem-level threats. These approaches should complement continuous integration and testing practices, forming a robust set of validation methods that can sustain accuracy, trust, and progress as AI becomes increasingly integrated into scientific computing.

### 3.2 Cross-disciplinary collaboration and AI for scientific software teams

The challenges of next-generation scientific computing demand teams that can operate across disciplinary boundaries, institutional cultures, and evolving technological landscapes. AI adds a new dimension to this complexity, both as a source of capability and as a potential collaborator. Research in this area must explore how scientific teams can effectively integrate AI systems into their workflows, while maintaining human creativity, trust, and scientific rigor. This includes developing new frameworks for trust, formalizing human-AI team roles, and applying lessons from prior collaborations in science, social science, and the private sector to guide effective engagement throughout the software lifecycle.

**A multilevel model of trust: Teams and technology.** Trust—foundational to effective collaboration—has been studied extensively as a concept; rich literature exists describing its presence in teams, differences between trust and distrust, and its application to technologies including those endowed with autonomy [69, 70, 71]. A multilevel, dynamic model of trust is needed to guide cross-disciplinary teams working in hybrid

HPC/AI environments. This includes constructs such as propensity to trust, perceived trustworthiness, trust in AI, trust in automation and automated systems, transparency, and reliance. Such a model must account for interactions across individuals, teams, and technologies, helping uncover opportunities to build and maintain trust, while also reducing the likelihood of cascading failures caused by misplaced trust or unaddressed uncertainty.

**Roles for AI: Teammate, trainer, tool.** As AI becomes increasingly embedded in scientific research, clarifying its roles within and across teams becomes essential. We must examine modes of human-AI teaming that augment human creativity and decision-making, especially in scientific code development and debugging. For example, AI-assisted collaboration tools could assist with translation between domain vocabularies (such as between scientific notation used by physicists and programming constructs familiar to research software engineers). To effectively integrate AI in teams and mitigate the likelihood of negative outcomes associated with the adoption of AI, research is needed to formalize and classify the various forms of taskwork and teamwork [60] in scientific computing. Doing so could enable, for example, the development of an adaptive taxonomy of roles and tasks for AI that is informed by the capabilities and limitations of the technology, the needs of the team, and the particulars of the domain. This configurable knowledge base could guide team design, risk management, and human-AI interaction strategies across the software lifecycle.

**Effective engagement within cross-disciplinary teams.** We must draw on insights from past scientific computing projects and other endeavors to foster environments that support extensive participation and cross-disciplinary success in next-generation computational science [61, 72, 73]. This includes updating institutional policies and incentives to recognize collaborative work [74] and embedding communication norms [75], governance structures, and knowledge-sharing practices that welcome contributions from a wide range of stakeholders throughout the entire software lifecycle—including funding acquisition; software design, development, and testing; developer and user (re)training; and use and maintenance of the software itself to conduct science. Efforts to codify best practices, define success metrics, and incentivize adoption are essential for building resilient, high-performing teams for the AI-augmented scientific future.

### 3.3 Pedagogy and workforce development in the age of AI

A perennial shortage of well-qualified scientists and engineers with the breadth of knowledge needed for today’s and tomorrow’s scientific computing underscores the urgent need for new pedagogical approaches, training programs, and workforce strategies. These efforts must support the development of cross-disciplinary skills, incentivize and promote ways to integrate different methodologies, ease onboarding into complex research projects, and enable continuous learning to keep pace with rapid technological evolution.

**Workforce development and human-AI collaboration models.** Research is urgently needed to understand how workforce development for scientific computing must evolve in AI-augmented environments. This includes empirical studies of how AI tools influence learning, skill development, and scientific reasoning among students and researchers at different career stages. Key questions include the following: Which fundamental skills remain essential when AI handles routine coding tasks? How can pedagogical approaches combine AI assistance with deep understanding of scientific principles? How should assessment methods evolve to evaluate comprehension in AI-assisted contexts? This research should also explore new models of human-AI collaboration in scientific software development, including optimal division of labor, interface design for effective teaming, and strategies for maintaining human expertise and critical thinking while leveraging AI’s computational advantages.

**Cross-disciplinary curriculum design.** Educational programs must integrate AI, software engineering, and scientific domain knowledge—from K-12 through graduate training. Identifying key competencies (e.g., trustworthy and fair AI, reproducibility, teamwork) is essential for preparing future software de-

velopers in scientific fields. Curriculum development must become more flexible and modular, enabling faster iteration and responsiveness to the changing landscape. This includes fostering partnerships across academia, national labs, and the private sector. Professional development programs should also address cross-disciplinary collaboration skills and provide transformational hands-on experiences that illustrate the impact and excitement of teamwork in scientific computing, while fostering collaborative environments that respect and adapt to different domain cultures.

**Innovative curricula and programs incorporating socio-technical co-design.** New instructional models are needed that emphasize project-based learning, especially through collaborations among the private sector, academia, and national laboratories. These models should extend beyond STEM domains to include the humanities and other relevant disciplines, ensuring a broader understanding of scientific and societal impact. The models should also include thoughtful integration of career pathways to support entry at multiple points, meeting people where they are, rather than requiring people to enter with a specific set of pre-defined prerequisites. In addition, we should capture lessons from successful efforts. For example, the collaboration model adopted by the Exascale Computing Project, in partnership with DOE leadership computing facilities, was highly successful in launching the exascale era by co-creating scientific applications, human capital, and a robust scientific software ecosystem. Not only did ECP bridge research and computing facilities in a co-design approach, ECP included early attempts at socio-technical co-design when it added talent development through the Sustainable Research Pathways [76] program and training pipelines through the Center for Scientific Collaboration and Community Engagement (CSCCE) [77].

**Community and workforce success metrics.** Clear, ambitious metrics are needed to measure progress in building the HPC+AI workforce. For example, a pilot program to advance the workforce for next-generation scientific computing could aim to increase the talent pool by at least 1,000 people over five years, averaging more than 200 new and upskilled contributors annually who are equipped to tackle emerging challenges in scientific software ecosystems. Well-defined metrics can help align national efforts, guide strategic investments, and ensure that workforce development keeps pace with evolving technical demands.

### 3.4 Required community actions

While technical research is a key driver of innovation in scientific computing, the workshop emphasized that progress in foundational scientific software also relies on sustained collaboration across institutions, disciplines, and sectors. Building trustworthy, capable, and sustainable software ecosystems requires more than technical breakthroughs—it calls for strategic coordination, shared infrastructure, supportive policies, and workforce investment. Various organizations, projects, and foundations—such as the Consortium for the Advancement of Scientific Software, High Performance Software Foundation, IDEAS project, NumFOCUS foundation, Research Software Alliance, Software Sustainability Institute, and US Research Software Engineering Organization—are making meaningful strides in shifting the cultural and structural foundations of scientific software [78]. Raising awareness of these efforts and contributing to them—whether directly or through complementary action—will benefit the broader scientific community. Much work remains to be done, however, particularly in advancing high-impact partnerships and addressing persistent gaps in infrastructure, recognition, and career paths.

**Strategic public-private partnerships.** Of the many strategies discussed, this emerged from the workshop as a central priority. Stronger collaboration among academic institutions, national laboratories, and the private sector is essential to the future of AI-enabled scientific software. These partnerships must go beyond short-term research goals to support co-designed solutions that align with long-term scientific needs and technological shifts. The private sector includes both commercial developers of computing technologies—such as hardware manufacturers, cloud service providers, and AI platform companies—and major

industrial users of scientific software. Companies across the automotive, aerospace, energy, manufacturing, and materials sectors (e.g., Ford, GE Research, Boeing, and others) depend on high-performance computing and simulation to support product design, optimization, and predictive modeling. Their engagement brings valuable application-driven insights, real-world performance requirements, and opportunities to accelerate the translation of foundational advances in computing into practical tools. Sustained coordination is needed to shape not only the technical requirements of software tools but also the practices and training ecosystems that support their adoption. Joint efforts might include cross-sector learning communities, affinity groups, shared training initiatives, and advisory partnerships that bring together HPC practitioners, AI developers, and domain scientists. Of particular importance is the need to align the rapidly evolving commercial AI and hardware landscapes with the long-term needs of the scientific community. These shared efforts play a vital role in maintaining national competitiveness, advancing research, and ensuring that science, engineering, and manufacturing progress together for long-term socioeconomic resilience.

**Community engagement and broad-access infrastructure.** As scientific software projects grow in scale and complexity, they require engagement models that foster trust, transparency, and shared ownership across teams and institutions. Building strong, multi-institutional collaborations depends on clear frameworks that support distributed leadership and far-reaching participation. Key strategies include implementing contributor agreements, project charters, and open governance models that reflect a wide range of stakeholder voices. Equally critical is sustained effort to ensure extensive, fair access to infrastructure and participation. This means funding and scaling infrastructure projects that lower barriers for under-resourced institutions and communities—such as through low-cost cloud access, peer mentoring networks, and national initiatives.

**Recognition, incentives, and long-term community investment.** The sustainability of scientific software ecosystems depends on valuing and supporting the full range of essential contributions, including development, documentation, user support, training, maintenance, and community stewardship. These roles must be reflected in hiring, promotion, and funding decisions and reinforced through long-term investment strategies. Stable, visible support for these roles strengthens the resilience of the overall ecosystem and builds the capacity needed for lasting progress.

**Responsible innovation guidelines and fair AI.** Advancing AI-driven scientific discovery without oversight regarding fairness risks opening a Pandora’s box of unintended consequences. Community-wide guidelines for the responsible integration of AI are essential and must be treated as a top priority. These guidelines should address algorithmic transparency, data governance, and reproducibility, while also establishing norms to prevent bias, protect users, and anticipate unintended consequences. A key part of this effort is ensuring that datasets are complete and representative of the full range of individuals and perspectives they aim to serve. Responsible innovation must remain a core design principle as both tools and teams evolve.

These policy and collaborative efforts are not secondary to technological innovation—they are foundational. By prioritizing strategic partnerships and reinforcing them with community-driven practices, institutional reforms, and long-term investment, while considering the long-term needs of both the nation’s research and industrial sectors, we can address the evolving demands of software- and AI-driven science. Ultimately, this integrated approach will foster ecosystems where technological breakthroughs and social progress advance hand in hand, reinforcing one another for a more innovative, responsible, and resilient future.

## 4 Conclusion and Next Steps

Scientific computing is entering a new era—driven by the integration of AI, the complexity of cross-disciplinary collaboration, and the growing demand for scalable, sustainable software ecosystems. As introduced in Figure 3, advancing next-generation scientific computing requires a socio-technical co-design

approach to fully harness community, software, and AI for cross-disciplinary team science. This means not only developing new software and hardware but also creating systems of practice—collaboration norms, training models, and governance strategies—that evolve with the technology. We launched a series of workshops to begin a dialogue in the computational science community and other related disciplines on how to confront and resolve these challenges. This report summarizes the outcomes of the first workshop in the series, representing the first step in a multiyear effort to understand how to co-design next-generation ecosystems for scientific computing. We envision ecosystems that are not only technologically advanced but also grounded in shared responsibility, fair practices, and a renewed commitment to bold, visionary science.

The workshop identified challenges and opportunities across three core domains: software infrastructure, cross-disciplinary collaboration, and pedagogy. Two crosscutting themes emerged: the transformative role of AI and the foundational importance of workforce and community development. Addressing these requires coordinated action on several fronts: advancing modular AI-integrated software; developing practices and tools that foster effective cross-disciplinary collaboration; and reimagining training for hybrid human-AI environments. Equally critical is aligning institutional policies, governance models, and community norms to ensure the long-term sustainability and growth of scientific software ecosystems and communities. Pilot projects can catalyze progress in these areas, with priorities evolving over time:

- **Near-term (1–2 years):** Launch pilot programs on hybrid AI/HPC software infrastructure as well as cross-disciplinary collaboration and pedagogy for AI-driven scientific computing, establish responsible AI guidelines, and prototype public–private partnerships.
- **Mid-term (3–5 years):** Explore scaling approaches for modular AI/HPC software ecosystems, launch workforce training curricula, expand community development, implement and evaluate several institutional policy reforms.
- **Long-term (5+ years):** Develop and evaluate globally networked ecosystems for AI-driven scientific computing, explore frameworks for community governance at scale, seek ways to integrate AI agents as active research collaborators.

Evolving next-generation scientific computing ecosystems, with many technical and social requirements, may seem daunting. However, the challenges also bring opportunities to learn from past successes and plan for agile approaches to grow communities and workforces poised to evolve alongside the changing landscape of HPC technologies and AI. Times of great change are marked by both enormous promise and significant uncertainty. Meeting this moment will require not just incremental improvements but bold, risk-tolerant thinking. Just as the dawn of quantum physics revolutionized our understanding of the universe, today’s scientific landscape demands a return to that spirit of curiosity, invention, and courage to chart new waters. Human ingenuity—coupled with emerging AI capabilities—must be placed at the center of scientific progress. By intentionally interweaving technical and social components in next-generation scientific computing, we can create feedback loops that accelerate both scientific discovery and real-world impact.

## Acknowledgments

This workshop was partially supported by the U.S. Department of Energy (DOE) Office of Science Distinguished Scientist Fellows Program. We especially thank our DOE contacts: Hal Finkel and David Rabson, DOE Office of Advanced Scientific Computing Research (ASCR).

We thank Katie Antypas (National Science Foundation, Office of Advanced Cyberinfrastructure), April Hanks (Team Libra) and Christina Mihaly Messina (SNL) for insightful contributions to workshop discussions on next-generation ecosystems in scientific computing, which helped shape the ideas conveyed in this report. We are grateful to Suzanne Parete-Koon (ORNL) for detailed feedback on the document; her suggestions improved the precision and perspective of the report. We thank Gail Piper for editing this manuscript.

## References

- [1] Anthropic. Anthropic Economic Index report: Uneven geographic and enterprise AI adoption, Sept 2025. <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>.
- [2] Bruce Hendrickson, Buddy Bland, Jackie Chen, Phil Colella, Eli Dart, Jack Dongarra, Thom Dunning, Ian Foster, Richard Gerber, Rachel Harken, et al. ASCR@ 40: Highlights and Impacts of ASCR's Programs. Technical report, US DOE Office of Science, 2020. doi:[10.2172/1631812](https://doi.org/10.2172/1631812).
- [3] William D. Gropp, Robert J. Harrison, et al. Future directions for NSF advanced computing infrastructure to support U.S. science and engineering in 2017–2020. National Academies Press, 2016. <http://www.nap.edu/catalog/21886/>.
- [4] David Keyes, Valerie Taylor, et al. National Science Foundation Advisory Committee on Cyber-Infrastructure, Task Force on Software for Science and Engineering, final report, 2011. [https://www.sci.utah.edu/publications/Key2011a/TaskForceReport\\_Software.pdf](https://www.sci.utah.edu/publications/Key2011a/TaskForceReport_Software.pdf).
- [5] U. Rüede, K. Willcox, L.C. McInnes, and H. De Sterck. Research and education in computational science and engineering. *SIAM Review*, 60(3):707–754, 2018. doi:[10.1137/16M1096840](https://doi.org/10.1137/16M1096840).
- [6] NSF Office of Advanced Cyberinfrastructure, CISE. Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century: **Blueprint for Cyberinfrastructure Learning and Workforce Development**, 2021.
- [7] Bruce Hendrickson et al. SIAM Task Force Report: The Future of Computational Science, 2025. URL <https://www.siam.org/media/ofajdru/siam-report-on-the-future-of-computational-science.pdf>.
- [8] Jonathan Carter, John Feddema, Doug Kothe, Rob Neely, Jason Puet, Rick Stevens, et al. Advanced Research Directions on AI for Science, Energy, and Security: Report on Summer 2022 Workshops, 2023. doi:[10.2172/1986455](https://doi.org/10.2172/1986455).
- [9] Michael A. Heroux, David E. Bernholdt, Lois Curfman McInnes, John R. Cary, Daniel S. Katz, Elaine M. Raybourn, and Damian Rouson. Basic Research Needs in The Science of Scientific Software Development and Use: Investment in Software is Investment in Science, 2023. Report of the DOE Advanced Scientific Computing Research Workshop, doi:[10.2172/1846009](https://doi.org/10.2172/1846009).
- [10] National Artificial Intelligence Research Resource (NAIRR) Pilot. National Artificial Intelligence Research Resource (NAIRR) Pilot, 2025. <https://nairrpilot.org>.
- [11] Michael E. Papka, Dhabaleswar K. Panda, Ilkay Altintas, Wahid Bhimji, Ewa Deelman, Murali Emani, Nicola Ferrier, Daniel S. Katz, Lois Curfman McInnes, Anita Nikolich, Feiyi Wang, and Jim Willenbring. Final Report of the 2024 NSF/DOE Workshop on NAIRR Software, 2025. US National Science Foundation Report, <https://events.cels.anl.gov/event/529>.
- [12] National Academies of Sciences, Engineering, and Medicine. Charting a Path in a Shifting Technical and Geopolitical Landscape: Post-Exascale Computing for the National Nuclear Security Administration, 2023. Washington, DC, The National Academies Press, doi:[10.17226/26916](https://doi.org/10.17226/26916).
- [13] Rick Stevens, Valerie Taylor, Jeff Nichols, Arthur Barney Maccabe, Katherine Yelick, and David Brown. AI for Science. DOE Report, OSTI, 2020. URL <https://doi.org/10.2172/1604756>.

[14] B. Chapman, H. Calandra, S. Crivelli, J. Dongarra, J. Hittinger, S. Lathrop, V. Sarkar, E. Stahlberg, J. Vetter, and D. Williams. DOE Advanced Scientific Advisory Committee (ASCAC): Workforce Subcommittee Letter. doi:[10.2172/1222711](https://doi.org/10.2172/1222711), 2014.

[15] Roscoe Giles et. al. Transitioning ASCR after ECP. [https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition\\_Report\\_202004-ASCAC.pdf](https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition_Report_202004-ASCAC.pdf), 2020.

[16] DOE Office of Critical and Emerging Technologies. Supercharging America's AI Workforce. <https://www.energy.gov/cet/supercharging-americas-ai-workforce>, 2024.

[17] Prasanna Balaprakash, Tina Brower-Thomas, Jennifer Gaudioso, Susan Gregurick, William D. Gropp, Arthur Maccabe, Irene Qualters, Mark E. Segal, Valerie Taylor, David Torres, and Stefan M. Wild. Response to the ASCAC charge to review the Computational Science Graduate Fellowship program. Report, U.S. Department of Energy, ASCAC, 2025. URL [https://science.osti.gov/-/media/ascr/ascac/pdf/reports/2025/CSGF-Subcommittee-Report\\_Final.pdf](https://science.osti.gov/-/media/ascr/ascac/pdf/reports/2025/CSGF-Subcommittee-Report_Final.pdf).

[18] Sudip Parikh et al. Vision for American Science and Technology (VAST): Unleashing American Potential, 2025. AAAS Task Force, [https://static1.squarespace.com/static/679a77b3c92f5d47013b8fd2/t/67bd2ca01a2e87002521ce78/1740450984121/VASTreport\\_FINAL\\_2025.pdf](https://static1.squarespace.com/static/679a77b3c92f5d47013b8fd2/t/67bd2ca01a2e87002521ce78/1740450984121/VASTreport_FINAL_2025.pdf).

[19] Saswata Hier-Majumder, Wahid Bhimji, Brandon Cook, Graham Heyes, Kalyan Kumaran, John MacAuley, Bronson Messer, Philip Roth, Jiachuan Tian, and Brice Videau. Post Exascale Software in the ASCR Facilities Ecosystem, December 2024. URL <https://www.osti.gov/biblio/2573945>.

[20] Anita Carleton, Mark H. Klein, John E. Robert, Erin Harper, Robert K Cunningham, Dionisio de Niz, John T. Foreman, John B. Goodenough, James D. Herbsleb, Ipek Ozkaya, Douglas Schmidt, and Forrest Shull. Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research & Development. Technical report, Carnegie Mellon University, Software Engineering Institute, 2021. URL <https://insights.sei.cmu.edu/library/architecting-the-future-of-software-engineering-a-national-agenda-for-software-engineering-research-development>.

[21] James Ahrens, Amber Boehlein, Rich Carlson, Joshua Elliot, Kjiersten Fagnan, Nicola Ferrier, Ian Foster, Lee Gimpel, John Shalf, and Dan Ratner. Envisioning Science in 2050, 2022. United States Department of Energy, Advanced Scientific Computing Research, doi:[10.2172/1871683](https://doi.org/10.2172/1871683).

[22] Edward O. Pyzer-Knapp, Jed W. Pitera, Peter W. J. Staar, Seiji Takeda, Teodoro Laino, Daniel P. Sanders, James Sexton, John R. Smith, and Alessandro Curioni. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. *npj Computational Materials*, 8, 2022. URL <https://doi.org/10.1038/s41524-022-00765-z>.

[23] Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624:570–578, 2023. URL <https://doi.org/10.1038/s41586-023-06792-0>.

[24] National Academy of Engineering and National Academies of Sciences, Engineering, and Medicine. Foundational Research Gaps and Future Directions for Digital Twins, 2024. Washington, DC, The National Academies Press, doi:[10.17226/26894](https://doi.org/10.17226/26894).

[25] Anthony Bilodeau, Albert Michaud-Gagnon, Julia Chabbert, Benoit Turcotte, J orn Heine, Audrey Durand, and Flavie Lavoie-Cardinal. Development of AI-assisted microscopy frameworks through

realistic simulation with pySTED. *Nature Machine Intelligence*, 6:1197–1215, 2024. URL <https://doi.org/10.1038/s42256-024-00903-w>.

[26] Philippa Spangenberg, Sebastian Bessler, Lars Widera, Jenny Bottek, Mathis Richter, Stephanie Thiebes, Devon Siemes, Sascha D. Krauß, Lukasz G. Migas, Siva Swapna Kasarla, Prasad Phapale, Jens Kleesiek, Dagmar Führer, Lars C. Moeller, Heike Heuer, Raf Van de Plas, Matthias Gunzer, Oliver Soehnlein, Jens Soltwisch, Olga Shevchuk, Klaus Dreibewerd, and Daniel R. Engel. msiFlow: Automated workflows for reproducible and scalable multimodal mass spectrometry imaging and microscopy data analysis. *Nature Communications*, 16, 2025. URL <https://doi.org/10.1038/s41467-024-55306-7>.

[27] DOE Exascale Computing Project (ECP). <https://wwwexascaleproject.org>, 2024.

[28] Francis Alexander, Ann Almgren, John Bell, Amitava Bhattacharjee, Jacqueline Chen, Phil Colella, David Daniel, Jack DeSlippe, Lori Diachin, Erik Draeger, Anshu Dubey, Thom Dunning, Thomas Evans, Ian Foster, Marianne Francois, Tim Germann, Mark Gordon, Salman Habib, Mahantesh Halappanavar, Steven Hamilton, William Hart, Zhenyu (Henry) Huang, Aimee Hungerford, Daniel Kasen, Paul R. C. Kent, Tzanio Kolev, Douglas B. Kothe, Andreas Kronfeld, Ye Luo, Paul Mackenzie, David McCalen, Bronson Messer, Sue Mniszewski, Chris Oehmen, Amedeo Perazzo, Danny Perez, David Richards, William J. Rider, Rob Rieben, Kenneth Roche, Andrew Siegel, Michael Sprague, Carl Steefel, Rick Stevens, Madhava Syamlal, Mark Taylor, John Turner, Jean-Luc Vay, Artur F. Voter, Theresa L. Windus, and Katherine Yelick. Exascale applications: skin in the game. *Phil. Trans. R. Soc.*, 378(2166):20190056, March 2020. doi:10.1098/rsta.2019.0056.

[29] E4S: An HPC-AI Software Ecosystem for Science. <https://e4s.io>, 2025.

[30] M.A. Heroux, S. Shende, L.C. McInnes, T. Gamblin, and J.M. Willenbring. Toward a cohesive AI and simulation software ecosystem for scientific innovation, 2024. RFI response to the Frontiers in AI for Science, Security, and Technology (FASST) Initiative, doi:[10.48550/arXiv.2411.09507](https://arxiv.org/abs/2411.09507).

[31] Lois Curfman McInnes, Michael A. Heroux, Erik W. Draeger, Andrew Siegel, Susan Coglan, and Katie Antypas. How community software ecosystems can unlock the potential of exascale computing. *Nature Computational Science*, 1:92–94, 2021. doi:[10.1038/s43588-021-00033-y](https://doi.org/10.1038/s43588-021-00033-y).

[32] Michael A Heroux, Lois Curfman McInnes, James Ahrens, Todd Gamblin, Timothy C Germann, Xiaoye Sherry Li, Kathryn Mohror, Todd Munson, Sameer Shende, Rajeev Thakur, Jeffrey Vetter, and James Willenbring. ECP libraries and tools: An overview. *The International Journal of High Performance Computing Applications*, 38(5):381–408, 2024. doi:[10.1177/10943420241271005](https://doi.org/10.1177/10943420241271005).

[33] Richard Gerber, Steven Gottlieb, Michael A. Heroux, and Lois Curfman McInnes. Transforming science through software: Improving while delivering 100x. *Computing in Science & Engineering*, 26(01):4–7, 2024. doi:[10.1109/MCSE.2024.3400462](https://doi.org/10.1109/MCSE.2024.3400462).

[34] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The AI Scientist: Towards fully automated open-ended scientific discovery. Technical Report 2408.06292, ArXiv, 2024. URL <https://arxiv.org/abs/2408.06292>.

[35] Kyle Swanson, Wesley Wu, Nash L. Bulaong, John E. Pak, and James Zou. The Virtual Lab: AI agents design new SARS-CoV-2 nanobodies with experimental validation. *bioRxiv*, 2024. doi:[10.1101/2024.11.11.623004](https://doi.org/10.1101/2024.11.11.623004).

[36] Ali Essam Ghareeb, Benjamin Chang, Ludovico Mitchener, Angela Yiu, Caralyn J. Szostkiewicz, Jon M. Laurent, Muhammed T. Razzak, Andrew D. White, Michaela M. Hinks, and Samuel G. Rodrigues. Robin: A multi-agent system for automating scientific discovery, 2025. URL <https://arxiv.org/abs/2505.13400>.

[37] J.S. Vetter et al. Extreme Heterogeneity 2018: Productive Computational Science in the Era of Extreme Heterogeneity, 2018. Report for the DOE ASCR Basic Research Needs Workshop on Extreme Heterogeneity, doi:[10.2172/1473756](https://doi.org/10.2172/1473756).

[38] Suren Byna, Stratos Idreos, Terry Jones, Kathryn Mohror, Rob Ross, and Florin Rusu. Report for the ASCR Workshop on the Management and Storage of Scientific Data, 2021. doi:[10.2172/1845707](https://doi.org/10.2172/1845707).

[39] Pavel Lougovski, Ojas D. Parekh, Joe Broz, Mark Byrd, Joseph C. Chapman, Yanne Chembo, Wibe A. de Jong, Eden Figueroa, Travis S. Humble, Jeffrey Larson, et al. Report for the ASCR Workshop on Basic Research Needs in Quantum Computing and Networking - 2023, 2024. doi:[10.2172/2430035](https://doi.org/10.2172/2430035).

[40] Douglas Kothe, Stephen Lee, and Irene Qualters. Exascale computing in the United States. *IEEE Computing in Science and Engineering*, 21(1):17–29, 2019. doi:[10.1109/MCSE.2018.2875366](https://doi.org/10.1109/MCSE.2018.2875366).

[41] DOE Scientific Discovery through Advanced Computing (SciDAC). <https://www.scidac.gov>, 2025.

[42] David Brown, James Hack, and Robert Voigt. The early years and evolution of the DOE Computational Science Graduate Fellowship Program. *Computing in Science & Engineering*, 23(6):9–15, 2021. doi:[10.1109/MCSE.2021.3120689](https://doi.org/10.1109/MCSE.2021.3120689).

[43] DOE Computational Science Graduate Fellowship (CSGF) Program. <https://www.krellinst.org/csgf>, 2025.

[44] NNSA Predictive Science Academic Alliance Program. <https://www.nnsa-ap.us/Programs/Predictive-Science-Academic-Alliance-Program>, 2025.

[45] NSF Science and Technology Centers. <https://www.nsf.gov/od/oiia/ia/stc>, 2025.

[46] NSF Engineering Research Centers. <https://www.nsf.gov/eng/engineering-research-centers>, 2025.

[47] Sustainable Horizons Institute. Sustainable Research Pathways. <https://shinstitute.org/sustainable-research-pathways>, 2025.

[48] James Ang, Andrew Chien, Simon David Hammond, Adolfy Hoisie, Ian Karlin, Scott Pakin, John Shalf, and Jeffrey S. Vetter. Reimagining Codesign for Advanced Scientific Computing: Report for the ASCR Workshop on Reimagining Codesign, 2022. US DOE Office of Science, doi:[10.2172/1822199](https://doi.org/10.2172/1822199).

[49] Timothy C. Germann. Co-design in the Exascale Computing Project. *The International Journal of High Performance Computing Applications*, 35(6):503–507, 2021. doi:[10.1177/10943420211059380](https://doi.org/10.1177/10943420211059380).

[50] National Research Council. Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification, 2012. Washington, DC: The National Academies Press, doi:[10.17226/13395](https://doi.org/10.17226/13395).

[51] National Academies of Sciences, Engineering, and Medicine. Reproducibility and Replicability in Science, 2019. Washington, DC: The National Academies Press, doi:[10.17226/25303](https://doi.org/10.17226/25303).

[52] Maya Gokhale, Ganesh Gopalakrishnan, Jackson Mayo, Santosh Nagarakatte, Cindy Rubio-González, and Stephen F. Siegel. Report of the DOE/NSF Workshop on Correctness in Scientific Computing, 2023. URL <https://arxiv.org/abs/2312.15640>.

[53] Nathan Baker, Frank Alexander, Timo Bremer, Aric Hagberg, Yannis Kevrekidis, Habib Najm, Manish Parashar, Abani Patra, James Sethian, Stefan M. Wild, and Karen Willcox. Workshop Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence. Report, U.S. Department of Energy, ASCR, 2019. URL <https://doi.org/10.2172/1478744>.

[54] Nick McGreivy and Ammar Hakim. Weak baselines and reporting biases lead to overoptimism in machine learning for fluid-related partial differential equations. *Nature Machine Intelligence*, 6:1256–1269, 2024. doi:10.1038/s42256-024-00897-5.

[55] Johannes Brandstetter. Envisioning better benchmarks for machine learning PDE solvers. *Nature Machine Intelligence*, 7:2–3, 2025. URL <https://doi.org/10.1038/s42256-024-00962-z>.

[56] J. Gregory Pauloski, Yadu Babuji, Ryan Chard, Mansi Sakarvadia, Kyle Chard, and Ian Foster. Empowering scientific workflows with federated agents, 2025. URL <https://arxiv.org/abs/2505.05428>.

[57] Stefan Wuchty, Benjamin F. Jones, and Brian Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, May 2007. doi:10.1126/science.1136099.

[58] H. M. Williams and N. J. Allen. Teams at work. In Julian Barling and Cary L. Cooper, editors, *The SAGE Handbook of Organizational Behavior: Volume 1 - Micro Approaches*, pages 124–140. SAGE Publications Ltd, July 2008. URL <https://doi.org/10.4135/9781849200448>.

[59] National Research Council. Enhancing the Effectiveness of Team Science, 2015. Washington, DC: The National Academies Press, doi:10.17226/19007.

[60] National Academies of Sciences, Engineering, and Medicine. The Science and Practice of Team Science, 2025. Washington, DC: The National Academies Press, doi:10.17226/29043.

[61] Stephen M. Fiore. Interdisciplinarity as teamwork: How the science of teams can inform team science. *Small Group Research*, 39(3):251–277, June 2008. doi:10.1177/1046496408317797.

[62] Jean Lave and Etienne Wenger. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, 1991.

[63] Etienne Wenger, Richard McDermott, and William Snyder. *Cultivating Communities of Practice: A Guide to Managing Knowledge*. Cambridge University Press, 2002.

[64] Marc Brooker and Ankush Desai. Systems correctness practices at Amazon Web Services. *Commun. ACM*, 68(6):38–42, June 2025. URL <https://doi.org/10.1145/3729175>.

[65] Jonathan Gorard and Ammar Hakim. Shock with confidence: Formal proofs of correctness for hyperbolic partial differential equation solvers, 2025. URL <https://arxiv.org/abs/2503.13877>.

[66] Sahana Chennabasappa, Cyrus Nikolaidis, Daniel Song, David Molnar, Stephanie Ding, Shengye Wan, Spencer Whitman, Lauren Deason, Nicholas Doucette, Abraham Montilla, Alekhya Gampa, Beto de Paola, Dominik Gabi, James Crnkovich, Jean-Christophe Testud, Kat He, Rashnil Chaturvedi, Wu Zhou, and Joshua Saxe. LlamaFirewall: An open source guardrail system for building secure AI agents, 2025. URL <https://arxiv.org/abs/2505.03574>.

[67] Sebastian Nicolajsen and Claus Brabrand. What is programming? *Commun. ACM*, 68(6):28–30, June 2025. URL <https://doi.org/10.1145/3713068>.

[68] Thomas Schmied, Jorg Bornschein, Jordi Grau-Moya, Markus Wulfmeier, and Razvan Pascanu. LLMs are greedy agents: Effects of RL fine-tuning on decision-making abilities, 2025. URL <https://arxiv.org/abs/2504.16078>.

[69] Jennifer Feitosa, Rebecca Grossman, William S. Kramer, and Eduardo Salas. Measuring team trust: A critical and meta-analytical review. *Journal of Organizational Behavior*, 41(5):479–501, June 2020. doi:[10.1002/job.2436](https://doi.org/10.1002/job.2436).

[70] Kurt T. Dirks and Bart De Jong. Trust within the workplace: A review of two waves of research and a glimpse of the third. *Annual Review of Organizational Psychology and Organizational Behavior*, 9(1):24–276, Jan. 2022. doi:[10.1146/annurev-orgpsych-012420-083025](https://doi.org/10.1146/annurev-orgpsych-012420-083025).

[71] P. A. Hancock, Theresa T. Kessler, Alexandra D. Kaplan, Kimberly Stowers, J. Christopher Brill, Deborah R. Billings, Kristin E. Schaefer, and James L. Szalma. How and why humans trust: A meta-analysis and elaborated model. *Frontiers in Psychology*, 14, March 2023. doi:[10.3389/fpsyg.2023.1081086](https://doi.org/10.3389/fpsyg.2023.1081086).

[72] Julia Koehler Leman, Brian D. Weitzner, P. Douglas Renfrew, Steven M. Lewis, Rocco Moretti, Andrew M. Watkins, et al. Better together: Elements of successful scientific software development in a distributed collaborative community. *PLOS Computational Biology*, 16(5), May 2020. doi:[10.1371/journal.pcbi.1007507](https://doi.org/10.1371/journal.pcbi.1007507).

[73] Lois Curfman McInnes, Michael A. Heroux, David E. Bernholdt, Anshu Dubey, Elsa Gonsiorowski, Rinku Gupta, Osni Marques, J. David Moulton, Hai Ah Nam, Boyana Norris, Elaine M. Raybourn, Jim Willenbring, Ann Almgren, Roscoe A. Bartlett, Kita Cranfill, Stephen Fickas, Don Frederick, William F. Godoy, Patricia A. Grubel, Rebecca Hartman-Baker, Axel Huebl, Rose Lynch, Addi Malviya-Thakur, Reed Milewicz, Mark C. Miller, Miranda R. Mundt, Erik Palmer, Suzanne Parete-Koon, Megan Phinney, Katherine Riley, David M. Rogers, Benjamin Sims, Deborah Stevens, and Gregory R. Watson. A cast of thousands: How the IDEAS productivity project has advanced software productivity and sustainability. *IEEE Computing in Science & Engineering*, 26(1):48–60, 2024. doi:[10.1109/MCSE.2024.3383799](https://doi.org/10.1109/MCSE.2024.3383799).

[74] James Howison and James D. Herbsleb. Incentives and integration in scientific software production. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pages 459–470, San Antonio Texas USA, 2013. ACM. doi:[10.1145/2441776.2441828](https://doi.org/10.1145/2441776.2441828).

[75] Reed Milewicz and Elaine Raybourn. Talk to me: A case study on coordinating expertise in large-scale scientific software projects. In *2018 IEEE 14th International Conference on e-Science (e-Science)*, pages 9–18, Oct. 2018. doi:[10.1109/eScience.2018.00010](https://doi.org/10.1109/eScience.2018.00010).

[76] Mary Ann Leung. Sustainable Horizons Institute. <https://shinstitute.org>, 2025.

[77] Center for Scientific Collaboration and Community Engagement (CSCCE). <https://www.cscce.org>, 2025.

[78] Daniel S. Katz et al. Community organizations: Changing the culture in which research software is developed and sustained. *IEEE CiSE*, 21:8–24, 2019. doi:[10.1109/MCSE.2018.2883051](https://doi.org/10.1109/MCSE.2018.2883051).

## A Glossary

This report aims to communicate across a broad community, including experts in high-performance computing, AI, computational science, social and cognitive science, team science, community and workforce development, and related topics. To support our goal of communicating clearly across disciplines, this glossary defines terms as used in this report.

- **AI for science:** The next generation of methods and scientific opportunities in computing, including the development and application of AI methods (e.g., machine learning, deep learning, statistical methods, data analytics, automated control, and related areas) to build models from data and to use these models alone or in conjunction with simulation and scalable computing to advance scientific research [13].
- **Cross-disciplinary:** Collaboration among multiple disciplines toward a shared objective [60].
- **Ecosystems for scientific computing:** Dynamic socio-technical systems made up of people, technologies, infrastructure, institutions, workflows, and cultural practices that collectively support the development and evolution of scientific computing. [page 2 of this report]:
- **Science of team science:** Research area in which scholars from various disciplines, such as psychology, organizational sciences, sociology, communication, and philosophy, contribute conceptually and empirically to understanding how science teams are organized and work together, how to best measure their effectiveness, and the implications of individual differences in team science [60].
- **Socio-technical co-design for scientific computing:** The intentional and integrated development of both technical components (e.g., software, AI, infrastructure) and social components (e.g., teams, institutions, practices, training) of ecosystems for scientific computing, ensuring that social and technical dimensions are addressed not in isolation but as tightly coupled elements of a unified, forward-looking strategy [page 4 of this report].
- **Taskwork:** Activities associated with achieving a team's goals [60].
- **Teamwork:** Interactions among team members that are essential for effective collaboration [60].
- **Team science:** Collaborative, interdependent research conducted by more than one individual [60].

## Acronyms

- **AI:** Artificial intelligence
- **HPC:** high-performance computing
- **CPU:** Central processing unit
- **GPU:** graphics processing unit
- **STEM:** science, technology, engineering, and mathematics
- **NSF:** National Science Foundation
- **DOE:** U.S. Department of Energy

## B Workshop Description

The 2025 Workshop on *Next-Generation Ecosystems for Scientific Computing: Harnessing Community, Software, and AI for Cross-Disciplinary Team Science* engaged over 40 cross-disciplinary experts in Chicago, IL from April 29 to May 1, 2025, to chart a path toward more powerful, sustainable, and collaborative scientific software ecosystems.

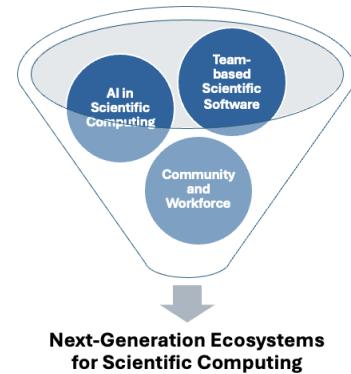
### Workshop Charge

The HPC community has long been a leader in advancing scientific discovery to new frontiers. High-quality scientific software—which encapsulates expertise across disciplines for use by others—is a primary means of sustained collaboration and scientific progress. Motivated by urgent issues raised in recent community reports on the state of scientific software development and AI for science, energy, and security, the workshop will bring together cross-disciplinary experts—in HPC, AI, computational science, applied math, computer science, research software engineering, cognitive and social sciences, and community development—to identify challenges, prioritize gaps, and explore opportunities that will shape next-generation ecosystems for scientific computing.

The workshop will follow a co-design methodology that intentionally weaves together topics in team-based scientific software, AI in scientific computing, and community/workforce development, thereby ensuring a holistic approach.

**Areas of emphasis.** The workshop’s goal is to assess and transform team-based scientific software, with emphasis on building ecosystems to address the needs of next-generation research in scientific computing, while advancing emerging AI technologies. We will integrate technical and community considerations throughout discussions on:

- **Software and next-generation science**
  - Exploring how new scientific frontiers and heterogeneous computing architectures require new approaches for software and workforce development.
  - Incorporating community-building strategies to ensure that new technical solutions reflect the perspectives and expertise of all stakeholders representing a broad swath of backgrounds.
- **Software ecosystems for AI in scientific computing**
  - Mapping pathways to create robust, scalable software ecosystems that catalyze AI-driven discoveries in HPC contexts and produce increased innovation.
  - Developing more robust AI tools and resources for scientific computing that address current gaps and include wide-ranging perspectives.
- **Team-based software and cross-disciplinary research**
  - Identifying roles, methodologies, and best practices for team-based scientific software.
  - Emphasizing community co-design, where software solutions evolve through iterative dialogue among scientists, developers, and end users.
- **AI for scientific software productivity and sustainability**



- Addressing how varied perspectives and AI can accelerate development, use, performance, and impact of scientific software.
- Embedding community-based feedback loops to align AI-driven tools with real-world needs and foster broader engagement.
- **Community and workforce development**
  - Accelerating strategies to cultivate future-generation R&D teams in the computing sciences.
  - Fostering collaborative environments that tap into the strengths of a broad workforce, bridging technical and community-building efforts.

## Workshop Objectives

We aim to shape a forward-looking plan for the future of team-based software in scientific computing, along with actionable strategies to bring it to life, all while cultivating a vibrant community. Key areas of focus—considering emerging AI technologies and the needs of cross-disciplinary research—will include:

- **Understanding scientific software practices:** Building a shared understanding of methodologies for team-based software in HPC, explicitly integrating technical and community perspectives.
- **Characterizing roles and success factors:** Identifying critical roles in scientific software teams and community-based success factors that sustain effective software development and user engagement.
- **Identifying challenges and opportunities:** Examining gaps (both technical and community-driven) that hinder team effectiveness and highlighting opportunities to address emerging research needs.
- **Envisioning the future:** Considering bold ideas for next-generation scientific software, emphasizing synergy between technical solutions and community dynamics.
- **Fostering community development:** Developing actionable steps to strengthen the workforce and build a vibrant software community prepared to meet urgent challenges in HPC and AI-driven scientific computing.

By intentionally blending technical and community factors, the workshop aims to shape a future where thriving, cross-disciplinary communities drive the next wave of scientific discovery, with high-quality scientific software as a keystone of sustained collaboration and scientific progress.

## Workshop Outcomes

Participants will map the landscape of team-based scientific software, characterize key roles and success factors, prioritize focus areas, and develop strategies to shape the future of scientific computing. These integrated insights—to be documented in a post-workshop report—will inform future efforts and may drive changes in perspectives, policies, and practices throughout the scientific computing community.

This workshop marks the first of a three-year series dedicated to strengthening team-based scientific software in an AI-driven future, with each year incorporating the co-design framework to ensure that both technical and community needs are addressed:

- **Year 1 emphasis:** Identifying and understanding challenges, gaps, and opportunities
- **Year 2 emphasis:** Developing and evaluating strategies and progress
- **Year 3 emphasis:** Coordinating ecosystem-wide advancements that meld technical solutions and community-building

By intentionally weaving community considerations into technical discussions—and vice versa—this three-year series will continuously expand the community base and help to advance next-generation scientific discovery.

## C Workshop Participants

### Organizing Committee

- Lois Curfman McInnes, Argonne National Laboratory (ANL), Senior Computational Scientist, Mathematics and Computer Science (MCS) Division
- Dorian Arnold, Emory University, Associate Professor, Department Computer Science
- Prasanna Balaprakash, Oak Ridge National Laboratory (ORNL), Director of AI Programs
- Mike Bernhardt, Team Libra, Founder and Chief Strategy Officer; former Director of Communication for the DOE Exascale Computing Project
- Beth Cerny, ANL, Head of Communications for the Argonne Leadership Computing Facility (ALCF)
- Anshu Dubey, ANL, Senior Computational Scientist, MCS Division
- Denice Ward Hood, University of Illinois Urbana-Champaign, Associate Professor, Department of Education Policy, Organization & Leadership
- Mary Ann Leung, Sustainable Horizons Institute, Founder and President
- Olivia B. Newton, University of Montana, Faculty, Department of Management Information Systems
- Stefan M. Wild, Lawrence Berkeley National Laboratory (LBNL), Director, Applied Mathematics and Computational Research Division

### Attendees

- Katie Antypas, National Science Foundation, Director of the Office of Advanced Cyberinfrastructure
- Tony Baylis, Lawrence Livermore National Laboratory (LLNL), Senior Manager
- David Bernholdt, ORNL, Distinguished R&D Staff Member, Computer Science and Mathematics (CSM) Division
- Chris Camañ, California Institute of Technology, Ph.D. Student, Applied and Computational Mathematics
- Hannah Cohoon, LBNL, User Experience Researcher, Scientific Data Division
- Charles Ferenbaugh, Los Alamos National Laboratory (LANL), R&D Scientist, Applied Computer Science
- Hal Finkel, DOE Office of Advanced Scientific Computing Research, Director of Computational Science Research and Partnerships
- Stephen M. Fiore, University of Central Florida, Professor, Director of Cognitive Sciences Laboratory,
- Sandra Gesing, US Research Software Engineers Association, Executive Director; Senior Researcher, San Diego Supercomputer Center
- Roscoe Giles, Boston University, Professor, Electrical and Computer Engineering, also Computing and Data Science
- Diego Gómez-Zará, University of Notre Dame, Assistant Professor, Computer Science & Engineering
- April Hanks, Team Libra, Director of Event Services
- James Howison, University of Texas at Austin, Associate Professor, School of Information
- Tanzima Islam, Texas State University, Assistant Professor, Department of Computer Science
- David Kepczynski, Ford Motor Company, Global Chief Digital Transformation
- Charles Lively, LBNL, Science Engagement Engineer and HPC Consultant, NERSC
- Vanessa López-Marrero, Brookhaven National Laboratory / Stonybrook University, Computational Scientist
- Harshitha Menon, LLNL, Computer Scientist, Center for Applied Scientific Computing,

- Bronson Messer, ORNL, Director of Science for OLCF
- Christina Mihaly Messina, SNL, HPC Developer & Metadata Manager, ORS Program
- Paul Messina, ANL, Argonne Associate and Distinguished Fellow, Computing, Environment and Life Sciences
- Marieme Ngom, ANL, Assistant Computer Scientist, ALCF
- Christopher Oehmen, Pacific Northwest National Laboratory, Senior Research Scientist and Group Leader, Biological Sciences Division
- Umesh Paliath, GE Aerospace Research, Technology Manager for Aerodynamics and CFD Methods
- Michael Papka, ANL, Deputy Associate Laboratory Director, Computing, Environment and Life Sciences; Director of ALCF; University of Illinois Chicago, Professor, Computer Science
- Irene Qualters, LANL, Associate Laboratory Director Emerita, Simulation and Computation
- David Rabson, Department of Energy, Advanced Scientific Computing Research, Physical Scientist
- Elaine M. Raybourn, Sandia National Laboratories (SNL), Social Scientist, Engineering & Software Research
- Katherine Riley, ANL, Director of Science for ALCF
- Paulina Rodriguez, The George Washington University, Ph.D. Student, Mechanical and Aerospace Engineering
- Damian Rouson, LBNL, Group Lead for Computer Languages and System Software, AMCR Division
- Michelle Schwalbe, National Academies of Sciences, Senior Director, Board on Mathematical Sciences and Analytics, Board on National Materials and Manufacturing
- Sudip Seal, ORNL, Group Lead for Systems and Decision Sciences, CSM Division
- Özge Sürer, Miami University, Assistant Professor, Department of Information Systems and Analytics
- Valerie Taylor, ANL, Director of the Mathematics and Computer Science (MCS) Division
- Jim Willenbring, SNL, Senior Member of R&D Technical Staff, Center for Computing Research
- Lou Woodley, Center for Scientific Collaboration and Community Engagement, Founder and Director
- Lingfei Wu, University of Pittsburgh, Assistant Professor, Department of Informatics and Networked Systems

## D Workshop Agenda

### Toward Next-Generation Ecosystems for Scientific Computing: Harnessing Community, Software, and AI for Cross-Disciplinary Team Science

#### Agenda

| Tuesday, April 29, 2025<br>Theme: <u>Challenges</u> in Next-generation Ecosystems for Scientific Computing |   |   |
|--|---|---|
| Time (CT)  | Topic   | Leads / Speakers / Panelists  |
| 8:00 am  | Registration  |   |
| 9:00 am  | Workshop commences  |   |
| 9:00 am  | Welcome to Argonne  | <b>Mike Papka</b> and <b>Valerie Taylor</b> (ANL)   |
| 9:10 am  | Workshop intro and logistics  | <b>Lois Curfman McInnes</b> (ANL) and <b>Mike Bernhardt</b> (Team Libra)  |
| 9:20 am  | Introductory remarks  | <b>Hal Finkel</b> , DOE Office of Advanced Scientific Computing Research, Director of Computational Science Research & Partnerships   |
| 9:30 am  | <b>Ecosystem dynamics:</b><br>Tackling challenges, unlocking opportunities, driving impact<br><br>Moderator: <b>Mike Bernhardt</b> (Team Libra) | <b>David Kepczynski</b> (Ford Motor Company, Global Chief Digital Transformation, Product Development)<br><b>Katie Antypas</b> (NSF, Director of the Office of Advanced Cyberinfrastructure)<br><b>Roscoe Giles</b> (Boston University, Professor, Electrical and Computer Engineering, also Computing and Data Science)  |
| 10:30 am   | Break   |   |
| 11:00 am   | <b>Grand challenges</b> in next-generation ecosystems for scientific computing<br><br>Moderator: <b>Stefan Wild</b> (LBNL)                      | <b>Katherine Riley</b> (ALCF, Director of Science)<br><b>Bronson Messer</b> (OLCF, Director of Science)<br><b>Umesh Paliath</b> (GE Global Research, Technology Manager for Aerodynamics and CFD Methods)<br><b>Stephen M. Fiore</b> (Univ of Central Florida, Professor, Director of Cognitive Sciences Laboratory)<br><b>Tony Baylis</b> (LLNL, Senior Manager)                                       |
| 12:00 pm   | Breakout instructions   |   |
| 12:15 pm   | Working lunch   | Participant intros  |
| 1:30 pm  | <b>Breakout session 1:</b><br><br>Grand challenges and gaps in next-generation ecosystems for scientific computing                              | <ul style="list-style-type: none"> <li>Software and next-generation science</li> <li>Software ecosystems for AI in scientific computing</li> <li>Team-based software and cross-disciplinary research</li> <li>AI for scientific software productivity and sustainability</li> <li>Community and workforce development</li> </ul> <b>Deliverable per group:</b> List of key challenges (w. dependencies) |
| 3:00 pm  | Break   |   |
| 3:30 pm  | <b>Lightning talks</b><br><br>Moderator: <b>Beth Cerny</b> (ANL)  | <b>Irene Qualters</b> (LANL) <b>Özge Sürer</b> (Miami Univ)<br><b>David Bernholdt</b> (ORNL) <b>Paulina Rodriguez</b> (GW Univ)<br><b>Vanessa López-Marrero</b> (BNL) <b>Marieme Ngom</b> (ANL)<br><b>Hannah Cohoon</b> (LBNL) <b>Chris Camañ</b> (Caltech)<br><b>Harshitha Menon</b> (LLNL) <b>Anshu Dubey</b> (ANL)<br><b>Jim Willenbring</b> (SNL)   |
| 4:00 pm  | <b>Cross-disciplinary discussions</b>   |   |
| 5:30 pm  | Adjourn   |   |

**Toward Next-Generation Ecosystems for Scientific Computing:  
Harnessing Community, Software, and AI for Cross-Disciplinary Team Science**

| Wednesday, April 30, 2025   |   |   |
|---|---|---|
| Theme: <u>Deeper Dive on Gaps and Opportunities for Next-generation Ecosystems for Scientific Computing</u> |   |   |
| Time (CT)   | Topic   | Leads / Speakers / Panelists  |
| 8:30 am   | Workshop day 2 commences: Intro to the day  |   |
| 8:45 am   | <b>Report-outs</b> from Day-1 breakouts: Challenges   | <b>Focus:</b> Key challenges (with dependencies among topics)<br>Moderator: <b>Anshu Dubey</b> (ANL)  |
| 9:15 am   | <b>Open forum:</b> Challenges in next-generation ecosystems for scientific computing  | Moderated discussion, emphasizing synergies among topics and co-design opportunities<br><b>Goal:</b> Overall list of key challenges (with dependencies)   |
| 9:45 am   | Break   |   |
| 10:15 am  | <b>Gaps and opportunities:</b><br>Next-generation ecosystems for scientific computing<br><br>Moderator: <b>Dorian Arnold</b> (Emory Univ) | <b>Prasanna Balaprakash</b> (ORNL, Director of AI Programs)<br><b>Paul Messina</b> (ANL, Argonne Associate, Founding Director of ECP, MCS, etc)<br><b>Elaine M. Raybourn</b> (SNL, Social Scientist, Engineering & Software Research)<br><b>Damian Rouson</b> (LBNL, Group Lead for Computer Languages and System Software)   |
| 11:15 am  | Breakout instructions   |   |
| 11:30 am  | Working lunch<br>Moderator: <b>Hannah Cohoon</b> (LBNL)   | Speaker: <b>James Howison</b> (University of Texas at Austin, Assoc Professor, School of Information)<br>Presentation: <b>Challenges and Pathways to Thriving Scientific Software Ecosystems</b>  |
| 1:00 pm   | Break   |   |
| 1:15 pm   | <b>Breakout session 2:</b><br>Gaps and opportunities in next-generation ecosystems  | <ul style="list-style-type: none"> <li>• Software and next-generation science</li> <li>• Software ecosystems for AI in scientific computing</li> <li>• Team-based software and cross-disciplinary research</li> <li>• AI for scientific software productivity &amp; sustainability</li> <li>• Community and workforce development</li> </ul>  |
| 2:45 pm   | Break   |   |
| 3:15 pm   | <b>Breakout session 2 (continued):</b><br>Gaps and opportunities in next-generation ecosystems  | <ul style="list-style-type: none"> <li>• Software and next-generation science</li> <li>• Software ecosystems for AI in scientific computing</li> <li>• Team-based software and cross-disciplinary research</li> <li>• AI for scientific software productivity &amp; sustainability</li> <li>• Community and workforce development</li> </ul> <p><b>Deliverable per group:</b> List of key gap and opportunities (with dependencies)</p> |
| 4:30 pm   | <b>Next-generation ecosystems for scientific computing: Project info</b>  | Speaker: <b>Olivia B. Newton</b> (Univ of Montana)  |
| 4:50 pm   | Summary of the day  |   |
| 5:00 pm   | Adjourn   |   |

**Toward Next-Generation Ecosystems for Scientific Computing:  
Harnessing Community, Software, and AI for Cross-Disciplinary Team Science**

| Thursday, May 1, 2025<br>Theme: <u>Next Steps: Building Next-generation Ecosystems for Scientific Computing</u> |  |   |
|---|--|---|
| Time (CT)   | Topic  | Leads / Speakers / Panelists  |
| 8:00 am   | Group discussions (optional) ... Crosscutting topics TBD   |   |
| 8:30 am   | Workshop Day 3 commences: Intro to the day   |   |
| 8:35 am   | Introductory remarks   | <b>David Rabson</b> , DOE Office of Advanced Scientific Computing Research, Physical Scientist  |
| 8:45 am   | <b>Report-outs</b> from Day-2 breakouts: Gaps and opportunities  | <b>Focus:</b> Key gaps and opportunities (with dependencies among topics)<br><br>Moderator: <b>Olivia B. Newton</b> (Univ of Montana)   |
| 9:15 am   | <b>Open forum:</b> Gaps and opportunities in next-generation ecosystems for scientific computing   | Moderated discussion, emphasizing synergies among topics and co-design opportunities<br><br><b>Goal:</b> Overall list of gaps and opportunities (with dependencies among topics)  |
| 9:45 am   | Break  |   |
| 10:00 am  | <b>Bright lights:</b> Some exemplar approaches for scientific computing ecosystems<br><br>Moderator: <b>Denice Ward Hood</b> (Univ of Illinois Urbana-Champaign) | <b>Dorian Arnold</b> (Emory Univ, Associate Professor, Department of Computer Science)<br><b>Lou Woodley</b> (Center for Scientific Collaboration and Community Engagement, Founder and Director)<br><b>Mary Ann Leung</b> (Sustainable Horizons Institute, Founder and President)  |
| 11:00 am  | Break  |   |
| 11:15 am  | <b>Breakout Session 3:</b><br><br><b>Next steps:</b> Building next-generation ecosystems for scientific computing  | <ul style="list-style-type: none"> <li>Software and next-generation science</li> <li>Software ecosystems for AI in scientific computing</li> <li>Team-based software and cross-disciplinary research</li> <li>AI for scientific software productivity and sustainability</li> <li>Community and workforce development</li> </ul> <p>For each subtopic in building next-generation ecosystems:</p> <ul style="list-style-type: none"> <li><b>What does success look like?</b></li> <li><b>How can we develop strategies to get there?</b></li> <li><b>How can we evaluate progress?</b></li> </ul> |
| 12:15 pm  | Break  |   |
| 12:30 pm  | Working lunch  | <ul style="list-style-type: none"> <li>Report-outs from Day 3 breakouts</li> <li>Open forum: Next steps: Building next-generation ecosystems for scientific computing</li> <li>Workshop outcomes, logistics, timeline</li> </ul>  |
| 2:00 pm   | Adjourn  |   |