The Role of Applied Mathematics in a New Era of Artificial Intelligence

SIAM AI Task Force Report | February 2026

Society for Industrial and Applied Mathematics

The Role of Applied Mathematics in a New Era of Artificial Intelligence

SIAM AI Task Force Report | February 2026

Executive Summary

Artificial intelligence is rapidly reshaping science, national security, healthcare, energy systems, agriculture, and the U.S. workforce. Federal agencies, universities, and industry are investing at unprecedented levels to deploy AI systems at scale and with urgency. Yet a critical gap is emerging: the rapid growth in AI investment is not being matched by investments in the applied mathematical sciences that underpin modern AI and are necessary for its reliable and sustained deployment and innovation.

- Without investments in applied mathematics, research AI will struggle to reach its full potential for high-priority high-consequence tasks that require enhanced trustworthiness and efficiency.
- With investments in applied mathematics and inclusion in interdisciplinary research, AI can become predictive, explainable, safe, and a driver of American competitiveness.

Applied mathematics is foundational infrastructure and a source for inventive new ideas. It develops the tools that allow AI systems to execute efficiently and obey physical laws, quantify uncertainty, generalize beyond training data, resist adversarial manipulation, and support decisions under real-world constraints. Without increased investment in this foundation, AI systems will not reach their full potential for predictive science tasks.

SIAM, representing more than 14,000 applied mathematicians and computational scientists, has prepared this report to articulate the reality and future vision of applied math's critical foundation to AI and to inform and engage with ongoing national discussions on AI strategy and priorities.

The current and envisioned AI capabilities described in this document cannot be realized by simply applying existing mathematical tools. New AI capabilities will build on decades of applied mathematics research and emerge by integrating mathematical rigor with modern AI systems. The question is how to continue the remarkable innovation story of AI, for which modern applied mathematics has been a major driving force. As AI reshapes scientific discovery, decision-making, and national infrastructure, it is simultaneously driving new and expanding applied mathematics research frontiers that must be advanced to sustain AI progress. Continued United States leadership in AI therefore depends on federal investment in applied mathematics research, alongside investments in AI deployment.

Why This Work Is Urgent

The United States is at a pivotal moment. Multi-billion-dollar federal AI investments, rapid commercial deployment, and university-wide AI initiatives are collectively setting the trajectory of the nation's scientific, economic, and security landscape for decades to come.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

At the same time:

- AI systems are being deployed faster than our ability to understand, verify, certify, or govern them.
- Federal agencies are launching large-scale AI programs with expectations for near-term impact, often without proportional investment in mathematical foundations needed to understand, evaluate and systematically improve AI systems.
- Universities are shifting hiring, curricula, and research infrastructure around AI, often without commensurate investment in applied mathematics, with lasting workforce consequences.
- Global industrial leadership is shifting from those with the most data to those with the mathematical insight to transform raw AI potential into reliable, efficient, scalable, and sovereign economic power.

Without strong mathematical foundations and continued mathematical innovation, these investments risk producing AI systems that are opaque, inefficient, or misaligned with national priorities. With mathematics at the core, they become engines of discovery, resilience, and competitiveness as well as a continued fountain of innovation.

The Core Message: What Applied Mathematics Delivers That AI Alone Cannot

Applied mathematics research has enabled critical characteristics of AI systems that are lost when AI is reduced to purely data-driven approaches, and it continues to generate the new methods required to make AI reliable, efficient, and trustworthy at scale. These characteristics and the relevant applied mathematics areas include:

- **Scientific validity**: enforcing physical laws, constraints, and first-principles reasoning
- **Efficiency and scalability**: surrogate models, reduced-order models, and fast algorithms for large-scale systems
- **Explainability and trust**: uncertainty quantification, causal reasoning, and formal interpretability
- **Safety and security**: provable bounds, adversarial robustness, verification, and validation
- **Generalization**: learning reliably from sparse, noisy, or incomplete data

These characteristics are not static; they depend on ongoing advances in applied mathematics research, particularly as AI systems grow more complex, data-hungry, and mission-critical. In short, applied mathematics investments will make AI predictive rather than merely correlational, and actionable rather than speculative, helping to quantify the uncertainty in the decisions made using AI systems. AI did not emerge apart from applied mathematics; rather, modern AI grew out of advances in applied mathematics research, and its continued success depends on continued innovation in the field. For progress to continue, this connection must remain healthy and vibrant.

Throughout this report, examples illustrate both what prior applied mathematics research has already enabled and what continued investment could make possible as AI systems are deployed at greater scale, complexity, and consequence.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

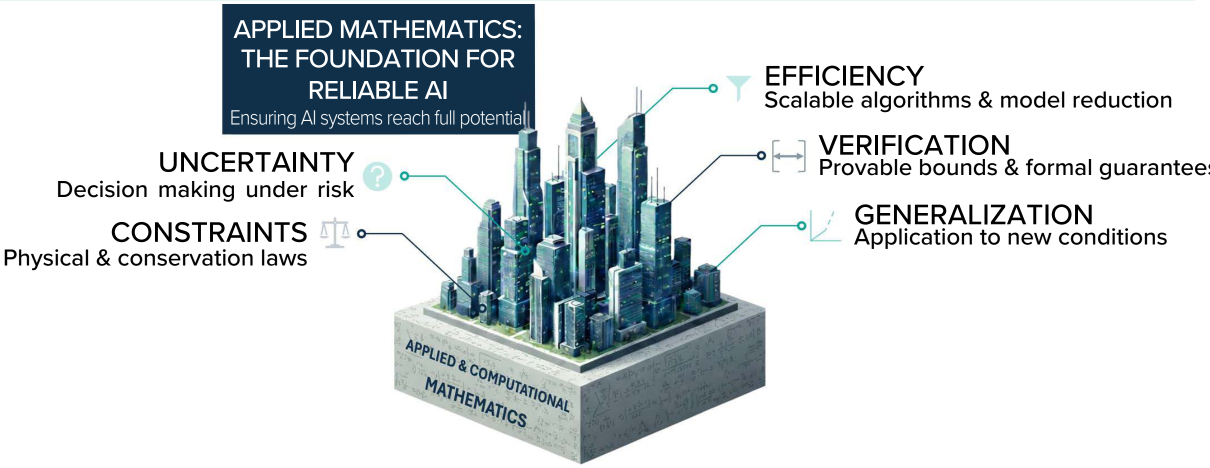

Applied Mathematics as National AI Infrastructure

Across domains, the same pattern emerges. Neither classical simulation nor modern AI alone is sufficient. The most successful approaches integrate:

- Data-driven learning
- Mechanistic modeling
- Mathematical constraints, verification techniques, and quantification of uncertainties
- Efficient and rigorous numerical algorithms as building blocks

This hybrid paradigm builds on the mathematical foundations of AI by integrating data-driven learning with mathematical structure and underlies emerging capabilities such as certified AI, autonomous experimentation, digital twins, and decision-support systems used in high-consequence settings. In each case, AI supplies scale and adaptability, while applied mathematics provides the rigorous methods required for prediction, verification, and control. Together, they enable systems that move beyond retrospective pattern analysis toward forward-looking, decision-relevant insight.

Digital twins illustrate this integration particularly clearly. For instance, in emerging healthcare applications, digital twins aim to represent an individual patient, not as a population-average statistical model, but as a personalized, predictive system combining real-time integration of medical data with mechanistic mathematical models to improve diagnostics and validate patient outcomes. AI plays a critical role in constructing, refining, and validating these models, but it is the mathematical foundations that give digital twins predictive and explanatory power: the ability to simulate "what-if" scenarios, quantify uncertainty, and adapt intervention and mitigation strategies as the patient's conditions change. Similar approaches are emerging across scientific discovery, energy systems, and national security, where digital twins enable decision-making under uncertainty rather than retrospective analysis. These capabilities are not achievable with data-driven AI alone; they depend on advances in applied mathematics that make AI predictive, verifiable, and actionable.

While the same foundational principles recur across domains, each application area places distinct demands on AI systems, highlighting different research frontiers for applied mathematics. The following five application areas illustrate how applied mathematics can transform AI from promise into capability: national security, healthcare, energy systems, agriculture and resilience, and physical sciences.

The Limits of Domain-Agnostic AI in High-Consequence Environments

As current AI systems are deployed in environments where failure carries irreversible consequences, their limitations become more visible. In such settings, systems must operate under uncertainty, incomplete information, and intentional disruption. They must predict rare events, distinguish signal from deception, and remain reliable when conditions change abruptly.

In adversarial and contested environments, today's AI systems that rely primarily on correlation lack the guarantees needed to predict behavior under stress. Statistical accuracy does not ensure robustness, verifiability, or safe operation when adversaries adapt or when data departs from historical patterns.

Rigorous mathematical methods provide the foundations required for future AI systems to function responsibly in these conditions. Such methods can represent physical constraints, quantify uncertainty, and enable verification. This will transform AI from a pattern-recognition tool into a decision-support capability suitable for mission-critical use.

The need for fundamentally new mathematical methods is most visible in national security applications, where the cost of error is measured not in performance metrics, but in strategic and operational risk. Meeting the demands of AI systems in these environments will require sustained applied mathematics research to develop new theory and algorithms for uncertainty quantification, adversarial robustness, optimization under constraints, and verification, as AI systems are deployed in contested, high-consequence environments where existing guarantees no longer suffice.

The national security example illustrates a broader point: wherever AI systems must support decisions under uncertainty, mathematics is the foundation that could make those systems trustworthy. The same mathematical principles that enable adversarial robustness and verification in national security contexts also underpin reliable AI deployment in civilian domains where safety and accountability matter. New mathematical techniques that can represent strategic and adversarial agents can enable early detection and robust response.

AI for National Security
High-consequence, adversarial systems demand AI tools that provide guarantees, not just predictions.

National security questions related to non-proliferation and the security of connected, interdependent infrastructures will require advances in AI that can handle low data and/or varied data environments. In particular, approaches that combine rapidly evolving foundation models with remote sensing data can be coupled with knowledge-informed system understanding to enable early detection of extremely weak and noisy chemical, biological, radiological, and nuclear (CBRN) signals. AI models, when abstracted from the mathematical foundations that originally enabled them, cannot solve these needle-in-haystack problems. Domain knowledge encompassing physical, environmental, and human context must be integrated using mathematical techniques to improve detection, deepen process understanding, and support inference about potential pathways and transport mechanisms for effective use of AI to enhance the warfighter. Understanding and responding to physical and cyber-security threats to our national critical infrastructures will also need detailed models of functioning infrastructure and their coevolution with the corresponding human systems. Dynamical representations of the functioning multi-scale and multi-networked systems using computational simulations are critical for detection, response, and counter-factual analysis. AI systems that rely solely on statistical pattern recognition, without embedded mathematical models or domain constraints, cannot reliably be used for counter-factual analysis, as they do not have the contextual information. In contrast, integrating AI with applied mathematical models makes them more reliable, interpretable, and suitable for high-consequence decision making.

In national security contexts, only continued mathematical innovation allows AI systems to move from analytical support to trusted strategic and operational decision-making.

Trust, Uncertainty, and Decision-Making Under Risk: From Predictive Accuracy to Clinical Trust

As AI systems increasingly inform real-world decisions, a fundamental limitation becomes clear: predictive accuracy alone is not sufficient for responsible use. In many high-stakes settings, decision-makers must understand not only what an AI system predicts, but how confident it is, when it may fail and how its recommendations should be weighed against competing risks.

AI for Healthcare
Clinical decision-making requires AI systems that communicate uncertainty, limitations, and rationale to clinicians and patients, not just diagnoses.

AI models trained to detect disease from medical images can achieve high average accuracy yet perform unevenly across hospitals with different patient populations or imaging protocols. Without uncertainty estimates, clinicians cannot determine when a prediction is reliable for an individual patient. Incorporating new applied mathematics tools into future AI systems will enable uncertainty-aware models that quantify confidence and flag out-of-distribution cases, allowing clinicians to integrate AI outputs safely into effective patient care decisions. Live, multi-scale digital twins of human body and organ systems will enable a new class of personalized predictive models that can lead to personalized and precision medicine. For example, new advances in treatment of diabetes using insulin pumps and blood glucose levels can advance the state-of-the-art of these closed-loop control systems. Dosing mechanisms for targeted cancer treatment/management and gut health treatment based on personalized diet and biome can lead to new therapies.

Mathematical algorithms extend AI's original diagnostic capabilities into clinically usable decision-support tools that can improve population health.

These challenges arise whenever decisions affect individual outcomes, involve incomplete or biased data, or carry ethical and legal responsibility. AI systems trained primarily to optimize performance metrics often struggle to provide transparent uncertainty estimates, to generalize across heterogeneous populations, or to explain how predictions change under different assumptions. When such systems are deployed without clear safeguards, they can undermine trust even when their average accuracy appears high.

In healthcare, applied mathematics provides the foundations needed to move from population-level prediction to individualized decision-support systems that can be trusted in practice. Mathematical frameworks enable explicit representation of uncertainty, enforce consistency with known constraints, and support reasoning about cause and effect rather than correlation alone. By making uncertainty visible, decisions auditable, and model behavior explainable to clinicians and patients, applied mathematics will allow AI outputs to be interpreted as inputs to human judgment rather than opaque conclusions.

These capabilities are especially critical in domains where decisions must be individualized, justified, and accountable, and where errors can directly affect human well-being. Continued investment in applied mathematics research is essential to extend and scale these capabilities so that AI systems can be deployed responsibly, with appropriate trust, oversight, and transparency.

These challenges make healthcare one of the clearest examples of why AI systems must be mathematically grounded to earn trust in real-world decision-making. Meeting these requirements depends not only on applying existing methods, but on continued applied mathematics research in uncertainty quantification, causal inference, interpretability, and decision-making under risk, as AI systems increasingly influence high-stakes clinical and public-health decisions.

Efficiency, Stability, and Control at National Scale: AI in Physically Constrained, Real-Time Systems

Many of the most consequential AI deployments will not perform isolated prediction tasks, but will rather be components of large, tightly coupled systems that must operate continuously and safely, for example the electrical grid. In these settings, AI systems must respond in real time to changing conditions, remain stable under perturbation, and respect hard physical and operational constraints. Small errors can propagate quickly, turning local inaccuracies into system-level failures.

Purely data-driven AI models, when decoupled from the mathematical structure that underpins physical systems, struggle in such environments because pattern recognition alone cannot guarantee stability, safety, or reliable behavior outside the conditions represented in training data. Systems governed by nonlinear dynamics, feedback loops, and conservation laws demand more than statistical approximation provided by AI. Without explicit constraints and verification, AI-driven decisions can amplify volatility rather than manage it.

Investments in formal methods will provide the structure that will allow AI systems to function reliably within these limits. Mathematical frameworks enable the representation of physical laws, the analysis of system dynamics, and the design of control strategies that remain robust under uncertainty. Techniques such as dynamical systems analysis, optimization under constraints, model reduction, and uncertainty quantification transform AI from a forecasting tool into a controllable component of complex systems.

These capabilities are essential wherever AI must operate at scale, in real time, and under strict safety requirements. In such contexts, efficiency is inseparable from stability, and innovation depends on guarantees rather than heuristics. New research into the interactions of AI and physical systems is needed to supply the foundations that make scalable, reliable AI possible.

The energy sector makes clear that scalable AI must be constrained, verified, and optimized, not merely trained. Achieving these capabilities requires sustained applied mathematics research to extend dynamical systems theory, optimization, uncertainty quantification, and control methods so they can provide guarantees for AI-driven, data-intensive energy systems operating at the national scale.

AI for Energy Systems

Real-time, physically constrained systems require AI that remains stable under stress, not just accurate on average.

Energy systems are undergoing a once-in-a-century transformation. Cascading failures in power grids need a detailed representation of the power grid as well as a first principles understanding of the power flows, control mechanisms, and hidden failure modes. Purely AI-driven models cannot predict cascades reliably as these cascades are a result of the coevolving control actions by relays and the nonlinear flow problems that result from these actions. Furthermore, as the power systems undergo significant transformation due to complex loads as well as generating units, we find ourselves in situations where data is not available to train the AI models under varying operating modes. Finally, new classes of threats, including adversarial as well as extreme weather events make contingency analysis even more challenging. In addition, cyber-security considerations are becoming increasingly relevant and need to be represented as a part of live digital twins. AI systems will need to be coupled with the digital twin of the underlying power grid to provide real-time assessment of potential cascades and automatic control schemes to reduce their impact. Foundational research in applied mathematics and co-design efforts are needed to design fast power flow solvers and dynamic transmission models, compute failure modes, develop novel importance sampling-based methods to identify rare events that can result in regional system collapse, and incorporate these methods into efficient and robust AI systems.

Mathematical models will allow AI to support real-time energy operations without increasing systemic risk.

Generalization, Sparse Data, and Multiscale Decision-Making: AI in Data-Poor, Environmentally Coupled Systems

Some of the most challenging environments for AI deployment are characterized by sparse, noisy, and heterogeneous data, coupled with dynamics that unfold across multiple spatial and temporal scales. In these settings, decisions must be made with incomplete information, long feedback loops, and interacting physical, biological, economic, and human factors. The central challenge is not prediction in isolation, but reliable generalization beyond the specific conditions represented in available data.

Purely data-driven AI systems often struggle in such environments because they are optimized to interpolate within historical datasets rather than extrapolate to new conditions. Models that perform well in controlled or localized settings may fail when applied across regions, seasons, or operating contexts. Without explicit structure, AI systems cannot reason about unseen scenarios, quantify risk under uncertainty, or support planning when conditions change.

Investments in applied mathematics research will provide the tools needed to address these limitations. Mathematical frameworks can enable AI systems to represent multiscale dynamics, integrate heterogeneous data sources, and propagate uncertainty across coupled subsystems. By embedding mechanistic structure and constraints into learning processes, mathematical innovation will allow AI to move beyond pattern recognition toward decision support, counterfactual analysis, and risk-aware planning. These capabilities are essential for systems in which outcomes depend on interactions across scales and where delayed or indirect effects matter as much as immediate predictions.

Such environments place especially high demands on generalization, robustness, and interpretability. In these cases, an understanding of mathematical structure will allow AI systems to remain useful when data are limited, conditions are changing, and decisions must account for long-term consequences. In agriculture and resilience planning, applied mathematics enables AI systems to generalize beyond historical data and remain useful as conditions change.

Agriculture highlights why generalization, not benchmark performance, is the true test of AI in real-world systems. Addressing these challenges requires ongoing applied mathematics research in multiscale modeling, generalization theory, uncertainty propagation, and decision-making under sparse data, as AI systems are pushed into complex, coupled environments where traditional assumptions no longer hold.

The agricultural case reinforces a central tenet of this report: AI systems deployed in complex, real-world environments must be grounded in mathematical models that capture structure, uncertainty, and scale. Without these foundations, AI becomes a powerful tool for resilience and sustainability. New mathematical methods to represent coupled human-natural systems can help address problems in food security while ensuring sustainable farming practices.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

AI for Agriculture and Resilience

Decisions across multiscale, data-sparse environments require AI systems that generalize beyond past observations.

AI models trained to predict crop yields from historical data may perform well in a single region but fail when applied to new climates, soil conditions, or extreme weather events. Without explicit structure, these models struggle to extrapolate beyond observed conditions. Advances in applied mathematics have enabled multiscale modeling and uncertainty propagation that allow AI systems to support risk-aware planning rather than brittle point predictions. Continued investment in applied mathematics research is needed to extend these capabilities to agricultural risk management challenges such as hurricanes, floods, droughts, and extreme weather events, where understanding propagation, compounding effects, and mitigation strategies is essential. These models can integrate complex topology, physical conditions, vegetation and urban landscapes to develop high-performance computing-oriented simulations that can then be coupled with AI models to develop real-time forecasts and mitigation measures. Mathematical models can be used to represent complex supply chains related to food production to markets and the risks to these integrated food systems due to extreme weather events. AI-based methods, when combined with mathematical models, can result in optimal water use for irrigation and result in early detection of pests.

New multi-scale mathematical models can enable AI driven robust decision systems when conditions change, and data are limited.

From Exploration to Predictive Discovery: AI as a Scientific Instrument

AI is increasingly used to accelerate discovery by exploring large design spaces, identifying patterns in complex data, and proposing new hypotheses. These capabilities have the potential to transform scientific and engineering workflows, enabling faster iteration and broader exploration than traditional methods alone. However, exploration is not the same as discovery. Scientific progress ultimately depends on results that are predictive, reproducible, and grounded in theory.

In scientific discovery, current AI systems that lack theoretical and mathematical grounding often struggle to meet standards of reproducibility and validity. In many scientific settings, data are sparse, costly, or hazardous to obtain, and the underlying systems are governed by physical laws, symmetries, the challenge of feedback between processes that are not necessarily known, and conservation principles that cannot be inferred reliably from data alone. Models that generate plausible outputs without enforcing these constraints may accelerate exploration, but they do not provide the guarantees required for scientific understanding or downstream use.

Innovations in applied mathematics will provide the foundations that allow AI to function as a scientific instrument rather than merely an exploratory tool. Mathematical frameworks enable the incorporation of physical laws, nonlinear feedback between processes, support causal reasoning, and provide rigorous methods for uncertainty quantification and validation. These tools will allow AI systems to distinguish signal from noise, to generalize beyond observed data, and to achieve high sample efficiency by abstracting fundamental structures from limited data, ensuring results that can be tested, reproduced, and built upon. Rather than replacing scientific reasoning, mathematics will ensure that AI augments it in a principled way.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

These capabilities are essential wherever AI is expected to contribute to fundamental understanding rather than superficial-level pattern discovery. In such contexts, speed without rigor is insufficient. Applied mathematics will ensure that accelerated discovery remains reliable, interpretable, and predictive, enabling AI to advance science rather than merely explore it. While accelerated discovery is often framed as merely speeding up discrete steps within the traditional scientific method, these mathematical foundations will enable a far more profound shift. We are moving beyond incremental speed gains toward a fundamental re-architecting of the discovery cycle itself. By unifying hypothesis generation, experimental design, and autonomous testing into an integrated loop, we can make breakthrough discoveries that cannot be entertained by a sequential approach (even with the fastest high-performance computing systems).

AI for Physical Sciences
Scientific discovery requires AI systems that are predictive and reproducible, not merely exploratory.

Generative AI models can rapidly propose candidate materials or experimental configurations, but without physical constraints they may produce unphysical or non-reproducible results. Examples where AI suggests unphysical solutions are well-documented. For example, using a trained AI model of engine data to find optimal operating settings resulted in a combustion engine that consumed, rather than produced CO_2, which is clearly nonsensical. Because this example involved a small number of parameters, these errors were easy to fix, but for larger models it may be harder to detect unphysical behavior. Applied mathematics enables physics-informed modeling, inverse problem formulation, and verification, allowing AI-generated hypotheses to be tested, refined, and trusted as part of the scientific discovery process.

Applied mathematics can extend AI's general capabilities into regimes where scientific discovery requires reliability, reproducibility, and predictive power.

In data-impoverished environments where traditional AI would fail, AI built with new applied mathematics will allow us to synthesize fragmented physical knowledge with sparse data. This fusion will enable the generation of breakthrough hypotheses that are fundamentally unreachable by purely data-driven or purely theory-based methods in isolation.

Realizing this vision depends on continued applied mathematics research to develop new theory and algorithms for inverse problems, uncertainty quantification, symbolic and hybrid modeling, and verification, that will enable AI to become embedded in the core processes of scientific discovery.

The physical sciences illustrate the full promise of mathematics-enabled AI: not only safer and more reliable systems, but faster and deeper discovery, unlocking a path to previously unreachable discoveries where scale alone is not enough. When AI is grounded in applied mathematics, it becomes a force multiplier for scientific leadership, technological innovation, and long-term economic growth.

Cross-Cutting Mathematical Capabilities

While the applications differ, they rely on a shared set of mathematical capabilities that cut across domains:

- **Uncertainty quantification and inverse problems**, especially in data-sparse or high-risk environments
- **Optimization and control**, including decision-making under constraints and uncertainty
- **Dynamical systems and multiscale methods**, enabling integration across time and spatial scales
- **Model reduction and surrogates**, providing better efficiency for and enabling real-time prediction and control
- **High-performance-computing-based numerical and discrete algorithms**, that leverage new computing hardware for faster and more accurate methods.
- **Verification, validation, and certification**, essential for deployment in regulated or adversarial settings
- **Fast multi-linear (tensor) and nonlinear algebra**, providing improved efficiency of AI operations

Sustained investment in these areas is critical to the success of AI across federal missions.

Implications for Federal Agencies

AI is not merely a new application area for applied mathematics; it changes the nature of applied mathematics research itself. Integrating learning systems with guarantees of stability, uncertainty, and correctness requires new theory, new algorithms, and new computational frameworks. Without advances in applied mathematics research, AI systems will remain brittle, difficult to certify, and costly to deploy, regardless of investment in data or compute power.

For federal agencies, recognizing applied mathematics as foundational to AI means:

- Explicitly including applied and computational mathematics in AI-focused efforts such as the Genesis Mission at the Department of Energy (DOE), AI Institutes and any new programs focused on AI for science at the National Science Foundation (NSF), Department of War (DOW) Multidisciplinary University Research Initiative, National Institute of Health AI initiatives, Advanced Research Projects Agency (ARPA) programs, and other team-based approaches to AI design and development.
- Supporting long-term and core mathematical research programs at NSF, DOE, and DOW that address foundational advances in mathematics and new mathematics tools that will seed later innovations specific to AI.
- Expanding partnerships like those between NSF, DOW, and NIH on digital twins and those in the DOE Science Discovery through Advanced Computing (SciDAC) partnership program that enable small teams to tackle specific challenges in combining mathematical innovations with science expertise into new AI systems.
- Massively expanding workforce opportunities that build mathematical skills for the future AI workforce and expanding the role of applied mathematicians in AI research efforts. Programs such as the NSF Research Traineeships and the DOE Computational Science Graduate Fellowship are effective existing vehicles that could be scaled. New initiatives should also be explored, including the AI Scholarship for Service program currently under consideration at NSF. In the health arena, there is a particularly acute need to strengthen connections and develop training that bridges applied mathematicians with the biomedical and health research community, equipping both groups with the skills necessary for effective collaboration.
- Treating verification, uncertainty, and reliability as key metrics for testing and validation through mechanisms such as the NIST Center for AI Standards and Innovation (CAISI) and the sandbox proposed in the Administration's AI Action Plan.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

Programs that integrate computational and applied mathematics from the outset, as illustrated across national security, healthcare, energy, agriculture, and the physical sciences, will deliver more reliable results, reduce downstream risk, and accelerate mission impact.

Implications for Universities

For university leaders, this moment presents both opportunity and responsibility. Durable AI programs require:

- Hiring strategies that strengthen applied mathematics alongside computer science
- Curricula that integrate modeling, development of efficient algorithms, formal verification, uncertainty quantification, and optimization into AI education
- Institutional support for interdisciplinary collaboration grounded in mathematical rigor and understanding

Universities that neglect these foundations risk producing graduates trained to deploy AI tools, but unable to build trustworthy AI systems.

Implications for Industry and Public-Private Partnerships

For industry, mathematically grounded AI reduces deployment risk, improves efficiency, and accelerates innovation. Strategic partnerships with applied mathematicians, including internships for graduate students and postdoctoral fellows, are essential to building AI systems that can be trusted in mission-critical environments, from energy infrastructure to healthcare delivery.

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 U.S.
Phone +1-215-382-9800 | Fax +1-215-386-7999 | www.siam.org | siam@siam.org

Contributors *(in alphabetical order)*

Alejandro Aceves, Southern Methodist University

Kevin Carlberg, University of Washington

Bert Debusschere, Sandia National Laboratory

Abba Gumel, University of Maryland

Aric Hagberg, Los Alamos National Laboratory

Lior Horesh, IBM

Vipin Kumar, University of Minnesota

Sven Leyffer, Argonne National Laboratory

Madhav Marathe, University of Virginia

Jonathan Mattingly, Duke University

Miriam Quintal, Lewis-Burke Associates LLC

Erin Raymond, University of Virginia

Karen Willcox, University of Texas - Austin

Carol Woodward, Lawrence Livermore National Laboratory