erVa

NSF Engineering Research Visioning Alliance

AI Engineering

A Strategic Research Framework to Benefit Society

Visioning Event Report

***AI Engineering*: A Strategic Research Framework to Benefit Society**

A visioning report of the Engineering Research Visioning Alliance

Report Finalized May 13, 2024

Based on proceedings from an ERVA event hosted by:

The University Financing Foundation, Inc.

This material is based upon work supported by the National Science Foundation (NSF) under award #2048419. Recommendations expressed in these materials are those of the authors and do not represent the views of the National Science Foundation.

The Engineering Research Visioning Alliance ([ERVA](https://ervacommunity.org)) is a neutral convener that helps identify and develop bold and transformative new engineering research directions, directly supporting the nation's ability to compete in a rapidly changing global economy. Funded by the National Science Foundation (NSF) Directorate for Engineering, ERVA is a diverse, inclusive, and engaged partnership that enables an array of voices to impact national engineering research priorities. The five-year initiative convenes, catalyzes, and empowers the engineering community to identify nascent opportunities and priorities for engineering-led innovative, high-impact, cross-domain research that addresses national, global, and societal needs. Learn more at ervacommunity.org.

Suggested citation: Engineering Research Visioning Alliance. 2024. *AI Engineering*: A Strategic Research Framework to Benefit Society. Columbia, SC: SSRN. dx.doi.org/10.2139/4778388.

Acknowledgments

ERVA is grateful to our partners at the National Science Foundation (NSF) Engineering Directorate for their ongoing engagement and support for this work, particularly Program Director Louise R. Howe and NSF Engineering Directorate leadership Susan Margulies, Don Millard, and Sohi Rastegar.

ERVA visioning events enable the engineering research community to identify nascent opportunities and priorities for engineering-led, innovative, high-impact research that addresses global and societal needs. Each event relies on the efforts of organizations and individuals who volunteer to lead, guide, and participate in its activities.

ERVA is grateful to its event co-host, The University Financing Foundation, for providing an excellent location for this event. We are especially grateful to ERVA co-PI Pramod Khargonekar, distinguished professor of electrical engineering and computer science and vice chancellor for research at the University of California, Irvine, who chaired the Thematic Task Force that developed the creative framework for this visioning event and served as primary report writer. Serving with him on the Thematic Task Force and helping to facilitate the event were Birgit "Bea" Braun, Dow; Mary "Missy" Cummings, George Mason University; Henry Jerez, Google; Vijay Kumar, University of Pennsylvania; Eric Miller, Tufts University; William H. Sanders, Carnegie Mellon University; and Pascal Van Hentenryck, Georgia Institute of Technology. ERVA PI Dorota Grejner-Brzezinska, The Ohio State University, lent expertise and oversight to the event. Susan Lang, director of the Center for the Study and Teaching of Writing at The Ohio State University, contributed to the report.

This visioning effort benefited from the assistance of annotators Ayomide Afolabi, Kennesaw State University; Ying Qian, University of Georgia; Yunji (Rosie) Kim, Georgia Institute of Technology; and Kayla Partee, Clark Atlanta University. ERVA and UIDP team members Josh Aebischer, Mike Brizek, Sandy Mau, and Rebecca Silveston supported event breakout facilitation, and Natoshia Goines supported event logistics.

Presentations via live video by Eric Topol, Scripps Research Translational Institute, and Rick Stevens, Argonne National Laboratory, offered topical context and informed the event.

Participants from diverse settings contributed to the event, representing large companies and startups, academic researchers, and government program officers. ERVA is grateful to all the workshop participants who contributed their time to create a valuable discussion and visioning report and to their organizations for the liberty to share their expertise for this effort.

Rebecca Silveston, ERVA executive director, led the event planning, program development, and execution. The event and report development were executed under the guidance of ERVA Principal Investigator Dorota Grejner-Brzezinska, The Ohio State University, and ERVA co-principal investigators Anthony Boccanfuso, UIDP; Charles Johnson-Bey, Booz Allen Hamilton; Pramod Khargonekar, University of California, Irvine; and Edl Schamiloglu, University of New Mexico. Development of the visioning session theme was informed by input from the ERVA Standing Council, Advisory Board, and NSF Directorate for Engineering collaborators.

Source: Canva abstract purple

Executive Summary

The strategic convergence of artificial intelligence (AI) and engineering, envisioned as *AI Engineering*, represents a generational opportunity to supercharge engineering for the benefit of society through enhancements to national competitiveness, national security, and overall economic growth. *AI Engineering* is a nascent field arising from this convergence and synthesis that will advance our nation's interests by leveraging the traditional strengths of engineering disciplines with breakthrough developments in the field. *AI Engineering* will be bidirectional and reciprocal: it evokes a future vision in which an engineering approach makes for better AI while AI makes for better-engineered systems. *AI Engineering* is based on the firm commitment of engineering processes and culture to ethics of safety, health, and public welfare and is a principal term used throughout this report.

Engineering for AI. Engineering disciplines will bring their domain knowledge, techniques, tools, and culture to the creation of new forms of explainable, trustworthy, and reliable AI-enabled cyber-physical systems. Knowledge, tools, and techniques from engineering will revolutionize the hardware on which AI is realized. Building on existing connections between AI and engineering fields such as signal processing, information theory, and control theory, the next generations of AI methods, models, and algorithms will be created. The most important benefits of these interdisciplinary fusions will come via the increased assurances of safety, reliability, and trustworthiness, as well as energy efficiency and environmental sustainability of AI-enabled cyber-physical systems.

AI for engineering. Increasingly capable AI tools can transform fundamental disciplines of engineering science. They will also transform major engineering endeavors of design, manufacturing, and infrastructure. These changes will impact the cost, performance, efficiency, customizability, and sustainability of engineered products and systems. They will also significantly enhance the productivity and capabilities of engineers across the full spectrum of the discipline: practicing engineers, engineering researchers, engineering educators, and engineering students.

AI Engineering has the potential to impact each of the 14 grand challenges articulated by the National Academy of Engineering.

The U.S. engineering enterprise is positioned to lead in the research and education necessary for the creation and development of *AI Engineering*, thereby enhancing U.S. leadership in AI and engineering technologies. Engineering researchers must assist with defining future AI systems through the evolution of existing and new systems, even as they employ existing AI systems to help drive the future of engineering.

AI Engineering has the potential to impact each of the [14 grand challenges] articulated by the National Academy of Engineering. To define *AI Engineering*, develop key strategies and initiatives, and identify innovative lines of research, a group of researchers, industry leaders, policymakers, and other stakeholders were brought together on Nov. 7- 8, 2023, at a visioning event convened by the [Engineering Research Visioning Alliance] (ERVA). During the two-day event, 28 participants generated and refined critical grand challenges at the intersection of engineering and AI that face engineering researchers now and in the next decade. These are listed below and explained in more detail in the full report.

Grand Challenges

01

Design, Manufacturing, and Operations

- Design safe, secure, reliable, and trustworthy AI systems.
- Transform manufacturing quality, efficiency, cost, and time-to-market through *AI Engineering*.
- Build and operate AI-engineered systems with cradle-to-grave state awareness.
- Overcome scaling challenges in engineering.

02

***AI Engineering* for Humans and Society**

- Construct engineered systems for safe, reliable, and productive human-AI team collaboration.
- Mitigate rare event consequences via AI.
- Incorporate ethics in all facets of *AI Engineering*.

03

National Initiatives for *AI Engineering*

- Enable collection, curation, and sharing of datasets to advance *AI Engineering*.
- Ensure equitable access to computational resources for *AI Engineering*.
- Develop engineering domain-specific foundation models.
- Establish dedicated research institutes for *AI Engineering*.
- Create new education and training programs for *AI Engineering*.


Source: Canva Human-AI

Taking Action

To realize these visions of progress in *AI Engineering*, it is incumbent upon leaders of government, universities, industry, and nonprofits, as well as professional engineering societies, to seize this generational opportunity, both in their sectors and in cross-sectoral collaborations.

Government

Breakthrough progress in both directions—engineering research for new AI, and AI for engineering research—should be enabled through targeted research and development programs at both state and federal levels, as well as computational infrastructure (e.g., GPU-based supercomputers available for research activities). Establishing new federal government programs (perhaps complemented by those at the state level or by private funders) to create multidisciplinary *AI Engineering* research institutes is both timely and urgent. Engineering leaders in government can also create in-house programs with a focus on *AI Engineering* to achieve their particular missions.

Industry

Industry leaders should create research and development programs to build capacity for breakthrough developments in *AI Engineering* in their specific industry sectors. They will need to bring cross-organizational and multidisciplinary focus on data, design, testing, and operations to both lead and take advantage of *AI Engineering*. This will be crucial for timely delivery of highly efficient, cost-optimized, and high-quality AI-engineered products and systems to market. Leveraging *AI Engineering* will dramatically enhance practicing engineers' productivity and capabilities.

Academia

Academic leaders must enable the full development of *AI Engineering* through their institutions' research, education, and service missions. Creation of multidisciplinary research programs and dedicated research institutes will enable their faculty, students, and staff to create the foundations of *AI Engineering*. They should also create education and training programs (degrees, certificates, continuing education, short courses) for their students and faculty as well as for their industry partners. They can leverage aspects of *AI Engineering* while empowering their local and regional innovation ecosystems.

Sector Convergence

In addition to work within their respective spheres, leaders from government, universities, industry, civil society, and nonprofits must collaborate to ensure successes. Strategic alignments among these sectors will energize collaborative efforts and will be essential to secure the financial, technological, organizational, and human resources needed to fully realize the *AI Engineering* vision. By garnering such resources, progress will be transformative. This sector convergence approach will facilitate a crucial element of the *AI Engineering* enterprise: the computing power and generation, collection, and curation of datasets for engineering-specific AI tools. For example, development of *AI Engineering* would benefit from an educational model in which government, industry and academia join together more tightly and at greater scale to support training for research engineers in this important, emerging discipline. Whether they have just received their bachelor's degree or have been working for decades, this next generation of students will pioneer *AI Engineering* through a training experience marked by crossing boundaries that had previously separated government, industry, and academic research campuses. Given intense global competition, a fully dedicated government-university-industry collaboration will put the U.S. engineering enterprise at a significant advantage. The immense vision and promise of *AI Engineering* dictate that we not miss this generational opportunity.

AI Engineering: An Introduction

AI Engineering is an emerging concept that organically combines the best of engineering and AI. In a reciprocal relationship, it leverages the idea that AI makes for better engineered systems while a systematic engineering approach makes for better AI. Our vision for *AI Engineering* is founded on two reciprocal, mutually supporting pillars.

The first comprises the ability to engineer improved technological systems that use AI. It is anticipated that AI will become an integral component of the full engineering cycle, from fundamental engineering sciences and engineering design to manufacturing, testing, operations, and recycling. These central engineering activities can be dramatically improved by systematically integrating current and future AI technologies. The widespread use of engineered industrial products and systems in all aspects of modern human civilization can make the potential impact of integration profound, particularly when the impact on process and systems transformation is considered.

The second pillar brings engineering domain knowledge, techniques, tools, and culture to the creation of new forms of explainable, trustworthy, and reliable AI-enabled cyber-physical systems. This pillar will build on and extend the frameworks from all relevant disciplines of engineering. Knowledge, tools, and techniques from engineering will revolutionize the hardware on which AI is realized. Building on existing connections between AI and engineering fields, such as signal processing, information theory, control theory, and traditional physics-based modeling and simulation techniques, the next generations of AI methods, models, and algorithms will be created. These advances will cover both directions in *AI Engineering*: prediction on the one hand, and decision-making or control on the other. The most important benefits of these interdisciplinary fusions will come via the increased assurances of safety, reliability, and trustworthiness, as well as energy efficiency and environmental sustainability of AI-enabled cyber-physical systems.

Semiconductor chips present a somewhat unique and interesting example of the confluence that *AI Engineering* will foster. On the one hand, devices such as CPUs, GPUs, custom AI processors, and programmable logic devices called field programmable gate arrays (FPGAs) are essential to modern AI systems. In turn, the design of current and future semiconductors using electronic design automation (EDA) tools is already influenced by the latest advances in AI. This virtuous cycle in semiconductor chips could serve as inspiration for all other domains of *AI Engineering* systems where AI and engineering mutually benefit one another. We also highlight the potential successors to semiconductor chips as the workhorse for AI computations, e.g., neuromorphic computing, and organoid intelligence. These will require leadership contributions from *AI Engineering*.

AI Engineering will be best positioned to complement humans in building a more prosperous society when developed in this bilateral manner. Engineering has a distinguished history of meeting human needs and aspirations. Engineered systems work within the constraints imposed by the laws of nature to meet human needs and desires. Using powerful design and synthesis techniques, engineers create safe and practical solutions to improve the human condition. The National Society of Professional Engineers' code of ethics makes the unequivocal commitment that "engineers shall hold paramount the safety, health, and welfare of the public." *AI Engineering* will build on these cultural foundations of engineering to create future technological systems that are productive, economical, reliable, and trustworthy, and that advance human flourishing.

To meet this potential and best serve the needs of society, *AI Engineering* must develop as a broad multidisciplinary field with strong theoretical foundations interlaced with compelling applications. These theoretical foundations will be enhanced by modeling across scales, modeling of human-AI teaming, and mathematical analysis, synthesis, and design methods and tools. *AI Engineering* will redefine how AI is approached from the practitioner's point of view, as well as from the academic point of view. *AI Engineering* curricula will educate a new generation of engineers who will work with experts from computer science, biology, medicine, physical sciences, social and behavioral sciences, law, business, and humanities to build reliable and trustworthy systems for use in industry, health, commerce, energy, climate, agriculture, and critical infrastructures. The educational landscape across fields will dramatically change by leveraging robust *AI Engineering* foundations. This change will be analogous to the widespread use of mathematics and computing in all engineering fields.

AI Engineering offers an optimistic vision of the future of humanity as the AI revolution gathers pace. For at least the last two centuries, the field of engineering has evolved to help human civilization solve major problems, create prosperity, and advance quality of life. *AI Engineering* can simultaneously leverage and shape the future development of AI to meet the challenges of the 21st century and beyond.

Grand challenge opportunities grouped in three broad categories are highlighted below. These are illustrative examples rather than an exhaustive description of the tremendous opportunities envisioned over the coming decade and beyond.

What is meant by AI in this report?

The term artificial intelligence (AI) has come to acquire many meanings since it was first coined at the legendary Dartmouth Summer Research Project in 1956. Indeed, the final report of the National Security Commission on Artificial Intelligence stated: “what Thomas Edison said of electricity encapsulates the AI future: ‘It is a field of fields … it holds the secrets which will reorganize the life of the world.’”

John McCarthy, a founding father of the field, defined AI as “the science and engineering of making intelligent machines, especially intelligent computer programs.”

The Organization for Economic Cooperation and Development defines the term thus:

> An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

In this report, we use the term AI to include the above definitions. As such, AI will include data science, various machine learning techniques and systems, natural language processing, reasoning, planning, and problem solving.

Grand Challenges

#1 Design, Manufacturing, and Operations

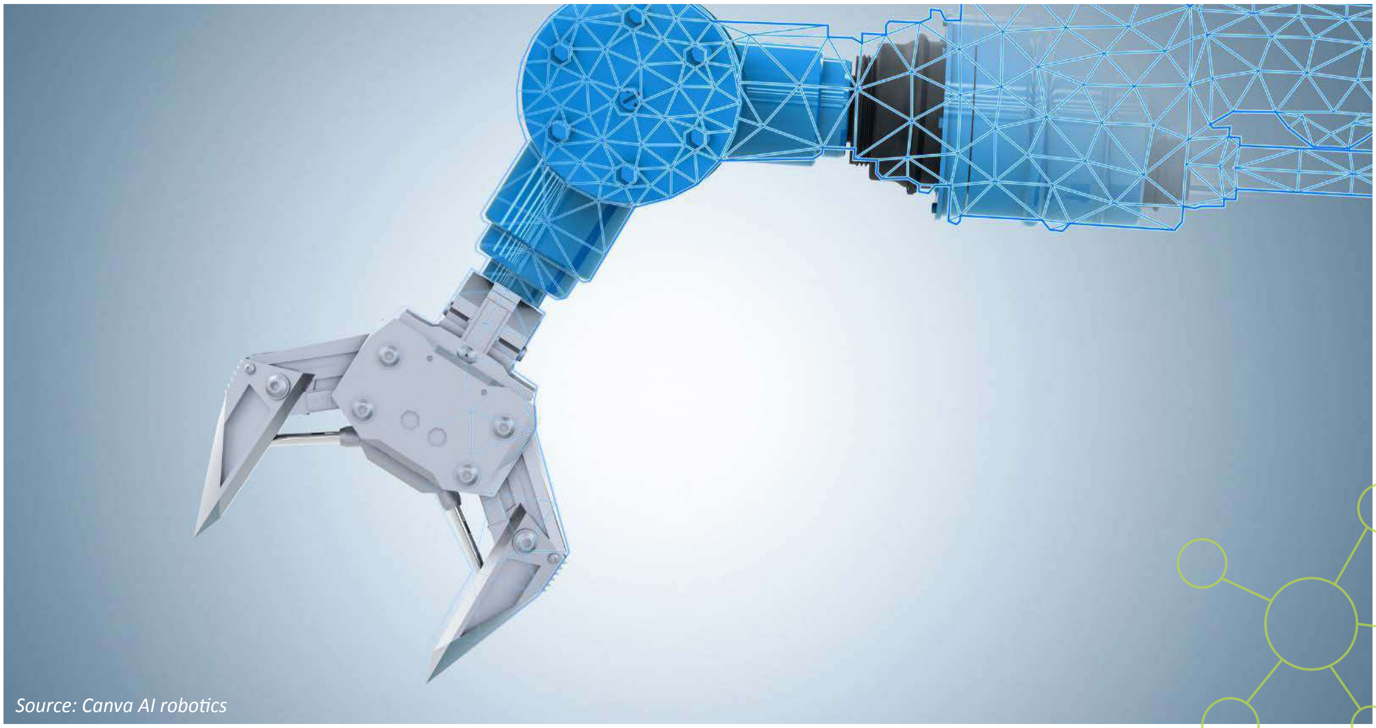

Source: Canva AI robotics

Current AI systems amalgamating data science, machine learning, generative AI, and their future development will dramatically transform the entire engineering design, manufacturing, and operations cycle. The potential benefits of such transformative changes include cost, efficiency, customizability, performance, cycle time, and sustainability. This transformation begins with accelerating advances in fundamental engineering science disciplines. These disciplines form the core basis of knowledge on which engineered products and systems are designed, manufactured, and operated. Examples of engineering science include fluid mechanics, thermodynamics, polymer science, and solid-state electronics. In some ways, impacts on engineering science fields will be like the impacts of AI on basic science fields such as chemistry, astronomy, and physics that are the focus of numerous recent reports. In essence, the widespread use of AI tools will accelerate the discovery phase of engineering science. The most visible examples to date of this accelerated discovery arise from the fields of materials science and engineering. Similar acceleration will occur in all fields of engineering sciences, although the rate of progress and specific opportunities will vary across fields.

AI tools, including generative AI, are already being investigated in certain subfields of engineering design, most notably in semiconductor chip design, leading to the next generation of EDA tools. A major opportunity exists for cross-sector learning in the integration of AI into engineering design.

Design Safe, Secure, Reliable, and Trustworthy AI Systems

AI safety is an emerging field that focuses on ensuring AI systems behave in a manner that is safe, reliable, and beneficial to humanity. It involves the study and application of principles, techniques, and methodologies to minimize the potential risks associated with AI while maximizing its benefits, especially in safety-critical settings.

AI safety has three distinct but complementary dimensions: (1) assuring a deployed AI system is safe and reliable, (2) using an AI system to monitor and improve the safety and reliability of a (potentially non-AI) system/platform, and (3) maximizing safety and trust in collaborative human-AI systems. It is critical to enumerate all three in any description of the science of AI safety. Closing the loop from identifying a new safety risk to resolving the root cause may take months to years in most complex systems, e.g., commercial aircraft and launch vehicles. This loop closure time can be drastically reduced when AI is introduced into the process, but how to certify AI-introduced system updates that respond to each newly identified safety risk remains an unanswered problem.

The transformative potential of this nascent and emerging engineering AI safety field lies in the ability of the engineering community to develop and deploy AI systems that are not only powerful and efficient, but also complement and enhance human activities. As AI systems become more prevalent and influential in society, ensuring their safety and reliability is increasingly critical. Ultimately, we will move from viewing AI systems as mere tools to AI systems as trusted team members that augment human abilities. A focus on engineering AI safety can help prevent harmful outcomes, mitigate risks, ensure that AI technologies are developed and used responsibly, and achieve their full potential.

Engineering researchers should lead research communities' efforts in the emerging field of AI safety because engineers specialize in taking technologies from fundamental concepts to deployment and implementation in society. An engineering AI safety program should include several facets. First, it should determine how to incorporate safety considerations into the design and development of AI systems, ensuring that they operate within defined parameters and do not pose undue risks. Second, engineers should develop AI safety principles and methods for monitoring and controlling AI systems, enabling timely intervention for unexpected outcomes. Finally, engineers should define and determine best practices for AI safety, demonstrating and promoting a culture of safety in the AI community.

Specific areas for high-impact research include:

Understanding Complex Systems: AI often interacts with complex social systems, where correlation versus causation might not be immediately clear. These interactions are multi-scale in time, space, and scope, with AI systems interacting with engineered systems and human beings. As such, safety, reliability, and trustworthiness imperatives increase in importance and complexity. Engineers must build partnerships with AI community members, stakeholders, and experts for deeper mutual understanding of these issues and the implications of making mistakes.

Data Management: AI systems require large amounts of high-quality and bias-free data for training and validation. Acquiring, preparing, and maintaining adequate bias-free data sets of high fidelity are key aspects of AI data management that must be formalized. Many AI systems also suffer from an explosion of data to train the models. Data cleansing, preparation, and secure technologies like homomorphic encryption are also crucial for efficient and safe handling of data in AI systems.

Secure, Trusted, and AI-optimized EDA: These tools are key to translating design ideas into physical implementations. Many steps in the EDA process require solving complex non-linear problems which are computationally demanding. AI has the potential to address these design optimization problems. A key problem for EDA companies is protecting the design data that trains EDA algorithms. Like proprietary manufacturing tool designs, a company's design data is key intellectual property; most companies would prohibit the use of design and related data to train general EDA algorithms that may help competitors equally. A new paradigm of API-driven EDA tooling is required to address this problem, allowing customers of such EDA tools to tailor the general EDA algorithms using their on-site custom AI models.

Testing and Certification of AI Systems: Because AI systems, particularly those based on machine learning, can be highly complex and non-deterministic, methods to test such systems and determine acceptable boundaries of performance are an open field of research. Many AI systems are designed to learn and adapt over time, so methods to design protocols to continuously monitor and test systems with embedded AI are also critical. Furthermore, the rapid pace of AI development means that any testing and certification may quickly become outdated, so determining the safest testing intervals, as well as necessary degrees of confidence in the results, must be explored.

Regulatory Considerations: AI applications often raise new regulatory (and ethical) questions. Engineers must navigate these issues while ensuring compliance with all relevant laws and regulations. To this end, engineering efforts are needed to define universally accepted standards and best practices for developing, testing, and deploying AI systems.

The goals of a research program in engineering AI safety and trust will be to define and achieve measurable improvements in safety and risk mitigation of systems with embedded AI. Progress in this research will result in building reliable systems that human operators and users can trust. Successes will include the development of techniques and testing for high-quality data as well as robust and verifiable AI, repeatable methods for value alignment, and frameworks for AI certification and regulation. In addition, integrating AI into existing systems and processes can be challenging. It requires careful planning and execution to ensure the AI system works seamlessly with other components.

Transform Manufacturing through Industrial *AI Engineering*

Problems encountered in complex engineering systems are often classified as visible and invisible spaces. Like an iceberg, the visible problem is only part of the entire issue, while the hidden problems are larger but unseen. Dealing with these problems can mean solving existing ones or heading off potential problems before they occur. *AI Engineering* will further accelerate the way engineers design products, optimize manufacturing operations, and create value for service systems. As AI becomes industrialized, the process will encompass data science and analytics, machine learning, cyber-physical systems, digital twins, and beyond. As more sensors and smart analytics software are integrated into networked industrial products and manufacturing systems, predictive technologies can further learn and autonomously optimize performance and productivity.

Newly emerging generative AI tools can enable gathering, understanding, and synthesizing “voice of the customer” quality feedback and user complaints, which today are labor-intensive processes. In engineering systems, decisions are often made using large knowledge models (including physical-based models, data-centric models, rule-based reasoning, and human experiences). One of the frequently used skill sets in product design and operations of complex engineering systems is exploring new design options, identifying root causes, and tracking solutions for a complex engineering system. This requires time-intensive efforts to recreate issues in lab environments so the appropriate solutions may be found.

Industrial AI requires a large amount of data with which to train AI algorithms. Companies that have invested in industrial digitalization (using information technology to increase margins, extend asset lifespans, and reduce environmental impact of operations) already have a steady volume of information streaming from their assets. Although most data sets are useful, they might not be useable for machine learning due to the lack of background and baseline information and brokenness of the datasets. According to a recent report from the National Academy of Engineering, the key issues to successfully develop and deploy AI for industrial applications include the lack of good quality data and lack of systematic approaches in developing and validating AI. Adoption is also hindered by AI trustworthiness and lack of standards.

AI-augmented analytics can help manufacturers boost productivity, accelerate innovation cycles, increase customer loyalty, and gain an advantage in the Industry 5.0 era. AI can be integrated into manufacturing applications in key areas, such as industrial robots that can adapt to their situations and cooperate with humans and other robots, semiconductors that can fine-tune their performance and minimize faults and errors, and production quality that can identify irregularities and apply corrective actions. Data-centric metrology systems are a critical area for smart semiconductor manufacturing, which can help yield improvement by overcoming inspection and metrology challenges through accelerated data-centric analytics.

Industrial *AI Engineering* research will advance a systematic discipline that focuses on developing, validating, and deploying various AI methods and machine learning algorithms for industrial applications with sustainable performance, scalability, and security. Combined with state-of-the-art sensing, communication, and big data analytics platforms, a systematic industrial AI methodology will allow integration of physical systems with computational models. It will also lead to an integrated learning platform to address the "3 D issues" – data, discipline, and domain – to strengthen interdisciplinary capabilities in diversified *AI Engineering* systems.

Build and Operate AI Engineered Systems with Cradle-to-Grave State Awareness

Cradle-to-Grave System State (CtGSS) awareness is a novel concept with the potential to revolutionize the way engineered systems are conceived, designed, manufactured, deployed, operated, and retired. Digital twins coupled with current and future AI technologies have the potential to help enable the development of the CtGSS concept. It is an engineering grand challenge that will only be fully realized through the confluence of AI and cyber-physical systems. When the technology needed to implement CtGSS is developed, it will provide detailed knowledge of all pertinent aspects of an engineered system throughout its lifecycle.

Depending on factors such as operational environment, aging, etc., these systems will need real-time adaptability of AI models (training and inference) on the edge of the network where systems are deployed. This necessitates highly efficient, real-time, reprogrammable computing systems on the edge of the network that are flexible, performant, power-efficient, and cost-effective.

Once the system is deployed with the AI, sensing, and controls that realize a CtGSS capability, the system will be continually monitored to assess its in-service performance and condition. It will possess the capability to implement appropriate controls to maximize service life and reliability, minimize operation and maintenance costs, reduce energy demand, and minimize environmental impact. AI will seek to make this in-service monitoring, assessment, and control process as autonomous as possible. Furthermore, it is envisioned that AI will provide continuous feedback to the material design and manufacturing processes to improve subsequent system designs.

Finally, CtGSS must provide for intelligent system retirement where end-of-life is predicted based on system state assessments and predicted degradation rates, anticipated future mission profiles, and sound engineering economic analyses. These assessments and predictions will be made more robust using AI. This process will seek to maximize recyclable components and safely contain and monitor hazardous non-recyclable components. Feedback will be provided to material design and manufacturing processes, as well as to the in-service monitoring, control, and assessment process, as part of a continual quality improvement cycle.

The CtGSS framework here complements findings surfaced in previous ERVA reports on distributed manufacturing and sustainable materials. It aligns with designing in system functionality at the material and manufacturing levels to minimize the use of raw materials and scrap (and their associated adverse environmental impact), the energy needed to develop and manufacture products, and the time from conception to finished product (with emphasis on small-lot or customized manufacturing). Similar to approaches envisioned for the design of sustainable materials, engineers must start planning for and building in "system intelligence" when a system is still a concept, prior to the operation and design phases. Implementation will require intelligent manufacturing process monitoring and control that will provide increased knowledge of the initial material, component, system, and system-of-system states. The AI needed to monitor and control the manufacturing processes will work in conjunction with new embedded sensing and controls capabilities, perhaps realized using multi-functional materials. CtGSS is applicable to all engineered systems, such as bioengineered artificial hearts, the next generation of inherently safe commercial nuclear reactors, satellites deployed for deep space exploration, and large-scale, one-of-a-kind scientific infrastructure systems such as particle accelerators.

The concept of CtGSS can only be realized through multidisciplinary technology development that will require expertise from almost all engineering and data science disciplines. Some researchers will develop the basic engineering hardware and software tools to implement this concept, while others will drive applications. Examples of disciplines and capabilities needed to realize this vision are listed below. A common theme across all these disciplines is that future advances will be dependent on the use of AI.

- **Engineering Mechanics (aerospace, mechanical, and civil):** Multi-scale, multi-physics modeling for damage initiation and evolution, modeling of transport phenomena;
- **Electrical Engineering:** Wireless telemetry, embedded systems, power management;
- **Chemical Engineering:** Energy storage, hydro-thermal waste destruction;
- **Computer Science:** Networking, AI, computing hardware architecture;
- **Information Science:** Large-scale data management, pattern recognition, signal processing;
- **Materials Engineering:** Multi-functional materials, material failure mechanisms, condensed matter physics;
- **Nuclear Engineering:** Transmutation of waste, thermal hydraulics;
- **Reliability Engineering:** Probabilistic risk assessment and reliability methods.

CtGSS will revolutionize the way complex "systems of systems" are designed, manufactured, maintained, and retired through the infusion of AI and cyber-physical systems concepts in all aspects of the engineering lifecycle process. These capabilities for decision making can be mediated, guided, and made beneficial for humans, communities, and institutions. This process will revolutionize the current conceive-design-build-deploy engineering process summarized in the well-known systems engineering V diagram below. It is to be noted that use of AI in engineered systems can not only significantly improve each of the boxes in the V diagram, but also reduce the time axis for such development, leading to decreased time-to-market, better quality, and cost effectiveness.

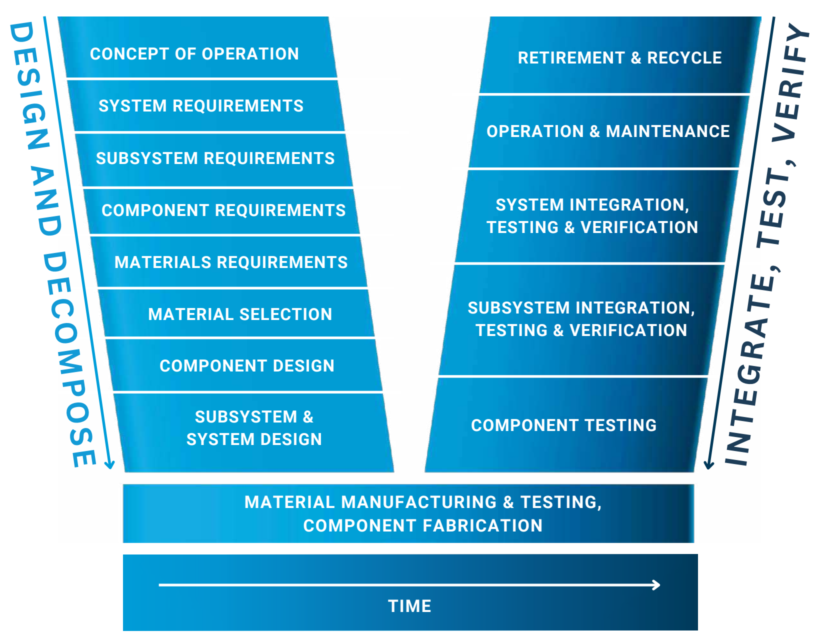

Source: Kossiakoff, A., Biemer, S.M., Seymour, S.J. and Flanigan, D.A., 2020. *Systems engineering principles and practice.*

Source: Canva AI monitoring

Overcome Scaling Challenges in Engineering

As we envision the future of *AI Engineering*, major cross-cutting and high-impact opportunity lies in scaling algorithms to very large problem instances and data sets. Today, engineering research, design, and analysis tools address these scaling problems through partitioning, layering, and integration. Three different engineering examples below illustrate where such scaling problems are especially pronounced.

Semiconductor scaling represents the first challenge. As Moore's Law slows, the industry has largely converged on multi-core and chiplet approaches to efficiently use silicon and create cost-effective and reliable computing solutions. Semiconductor chips implementing processors and/or accelerators are ubiquitous. As these chips become denser and more complex, the space of parameters that must be searched, implicitly or explicitly, to obtain efficient architectures that optimize performance, power, and area (PPA) becomes immense. It is often human intuition guided by trial-and-error experimentation that determines future computer architectures. However, these architectures are not necessarily optimal for computing problems that are very rapidly evolving (especially with the advent of AI).

Emergent system behavior due to evolution of component properties poses the second challenge. The combination and integration of very large numbers of even simple components may lead to the emergence of complex behaviors that cannot be easily predicted from that of the individual components. Different materials may exhibit new properties when integrated into a component, and components may exhibit new (unknown) properties and behaviors when integrated into a system or engineered product. Environmental conditions also affect the behaviors of these systems. This necessitates modeling, sensing, and prediction to avoid catastrophic events.

Software solvers are used by many engineering domains, especially EDA, to solve NP-hard problems, such as integer linear programming, Boolean satisfiability solvers, automatic test pattern generation solvers, or non-linear optimization techniques like conjugate gradient, interior point methods, simulated annealing, genetic algorithms, or Lagrangian optimization. These techniques are computer-intensive and have many built-in heuristics that require constant tuning and improvement. Given their importance in many domains, these solvers must become more efficient by orders of magnitude without loss of accuracy or the quality of results.

AI Engineering research to systematically address these scaling challenges will have a transformative impact. Below are preliminary suggestions for research directions that could be explored.

Semiconductor Scaling

AI may enable the next breakthrough in processor and accelerator technology (and the associated programming models) by analyzing large amounts of data to create accurate predictive models of key workloads, reducing the need for time-consuming simulations. Such AI-driven models and chip design methodology can help in significantly reducing the search-space for the parameters used to guide next-generation computer architecture research.

This research will interest architects and researchers in CPU, GPU, and custom accelerators/FPGAs, as well as those in related fields, such as manufacturing, packaging, testing, 3-D-integration, and compiling. AI itself is driving a huge demand for more efficient computing, and breakthrough solutions are required. Deep collaboration is essential between industry and academia, hardware engineers and researchers, and software developers and researchers. Reliable data insights and knowledge from prior efforts and workloads are needed to advance the effectiveness of AI-driven approaches for computer architecture and associated software programming. Research in these areas will lead to more efficient and cost-effective computer architecture and faster time-to-market to address increasing computing demands.

Emergent System Behavior Due to Evolution of Component Properties

AI can support consistent and self-adaptive co-design across space and time scales of engineered products spanning materials, components, and systems. AI can be used to predict (unknown) behaviors when materials are integrated into a component, when components are integrated into a system, or when systems are used in varying environmental conditions.

All system manufacturers, component designers, and materials scientists will be interested in this research. New fundamental research is needed in predicting system-level and component-level behaviors and properties in various environmental conditions of system operation with emerging self-learning systems.

Software Solvers

AI-driven solvers can help with dramatically reducing the search space for these solvers and even make them self-adaptive to the problem they are solving. This would be a new fundamental frontier for engineering research, with potential benefits for many domains. AI-driven software solvers do not necessarily change the fundamentals of the algorithms themselves but can markedly improve decision making during execution. This guidance can be domain-specific, problem-specific, and self-adaptive, and advances can have great impacts on many engineering domains.

Software solvers are used in every industry, particularly EDA, to solve various optimization problems. AI could invigorate new fundamental research directions for software solvers for computationally hard constrained-optimization problems prevalent in almost every engineering domain. A key result would be for these next-generation solvers to scale efficiently to problem sizes that are a few orders of magnitude larger than the problems dealt with today, as well as to produce better quality of results with an order of magnitude improvement in computational effort.

Several common threads can be discerned in the deployment of AI tools to address the scaling problems articulated above. AI tools are generally deployed to accelerate computationally burdensome elements of existing engineering workflows; how these tools might lead to radically new workflows remains an important topic of exploration. AI tools require large amounts of training data, which in turn requires extremely large amounts of computing time to generate AI models, especially when using supervised learning approaches. The amount of training data required to provide reliable predictions—and the ability of these tools to generalize beyond the data sets used to train them—are extremely poorly understood at present. Similarly, the use of synthetic data to leverage limited empirical data must be approached with caution as it is possible to introduce bias and unintentional statistical properties that may be difficult to detect. Domain-specific AI model development with self-adaptive behaviors needs further research and is critical to enable large-scale AI model training and real-time inference.

#2 *AI Engineering* for Humans and Society

Source: Canva AI general

Construct Engineered Systems for Safe, Reliable, and Productive Human-AI Team Collaboration

There is a long-running historical debate and competing visions between proponents of artificial intelligence and intelligence augmentation. The *AI Engineering* perspective is that of intelligence augmentation, which aims to empower human beings rather than replace them. At its core, engineering serves the needs of diverse groups of people who seek practical and ethical solutions to complex societal problems. AI certainly holds great promise to vastly improve all aspects of this mission. To do so most effectively will require careful consideration of how best to incorporate these computational tools into human engineering teams. This integration will require new engineering research that converges *AI Engineering* with human-centered design, social and behavioral sciences, humanities, and the arts.

At a foundational level, understanding and ultimately optimizing the interactions of heterogeneous groups of people and AI agents across multiple scales of space and time stands at the intersection of engineering, cognitive science, and social science. Human behavior is characterized by enormous variability, uncertainty, and stochasticity. Respect and concern for human autonomy is a core ethical principle. These facets of human behavior and respect for autonomy make human-AI collaboration extremely challenging. Such collaboration goes well beyond human-AI interaction. It will require major scientific and engineering advances in mutual understanding of goals, sharing of tasks, and tracking of progress. As such, this area is likely to remain a grand challenge for decades to come.

From smartphones, watches, and rings to video and audio sources, advances in our ability to collect “personalized” data *in situ*, as opposed to a controlled lab, are yielding an explosion of data streams, all of which have the potential to provide substantial insight into how engineers collaborate *and* how AI tools can help facilitate this work. Opportunities in this space will increase with advances made by electrical, chemical, material, mechanical, and biomedical engineers, as well as computer and data scientists. To date, some of these include antenna arrays to determine heart and respiration rates as well as flexible, wearable sensors capable of providing long-term, persistent, and unobtrusive monitoring of human bio-signals.

While it is possible to imagine the potential of such data for human-AI teaming, turning such data into useful teaming strategies is a major challenge and research opportunity. The data sets will be highly heterogeneous; some will be immense, while others will lack sufficient numbers in the sample to capture the full range of variability encountered. There will also be data missing from these sets in a myriad of interesting and frustrating ways. Ultimately, it will be necessary to develop algorithms that can leverage such data for instructing AI to assist humans in reaching goals specified by human agents.

Engineering researchers with expertise in dynamical systems theory, control, signal processing, and optimization will collaborate with practicing engineering teams, ethicists, cognitive scientists, social scientists, and computer and data scientists to define the questions to be answered with the data. These groups will develop fundamentally new approaches for constructing and guiding teams of humans and AI agents seeking to solve engineering problems collaboratively. As human-AI collaboration matures, both the cognitive and affective dimensions of trust must be considered to achieve high effectiveness.

Mitigate Rare Event Consequences via AI

A robustly engineered system accounts for and mitigates rare events of high consequence, such as engineered features that avoid bridge collapse, pharmaceutical side effects, chemical plant failures, security flaws, and seismic-induced damage to buildings. Engineers rely on domain knowledge, risk analysis, disaster forecasting, and other safety factors and considerations to bolster a system against rare events. Considering the advantages AI brings to many engineering applications, incorporating AI into the current process to create robust systems is compelling. However, current AI methods do not perform well in situations that are data-starved or data-absent. Balanced training sets that comprise information on rare events are inherently difficult or impossible to obtain. Thus, designing AI methods that are not blind to unforeseen or unavoidably complex outcomes throughout the service life of an engineered system requires significant advancements.

Source: Canva AI traffic

Past AI developers benefitted from data-rich applications and user bases with a high tolerance for failure; for example, machine learning-based language translation algorithms resulting in confusing, non-sensical, or humorous errors. Conversely, early-stage efforts in self-driving vehicles, where mistakes can cause fatalities, have suffered from slower adoption. Research teams in this space have built out test tracks, driving simulations, and/or massive fleet-wide data aggregation mechanisms to improve performance. These efforts aim to grow awareness of a vast problem space comprised of an unknown number of uncommon or unexpected scenarios that can be exceedingly important. The barriers to incorporating AI into larger-scale engineering systems will be greater without historical or growing relevant datasets, risk profiles that account for loss-of-life and costly investments, and unique applications in which knowledge transfer from other domains is not straightforward. Fortunately, the engineering community can draw upon established techniques from the fields of statistics, control theory, and disaster forecasting to add AI to the toolset that mitigates risk against high-consequence, low-frequency events characteristic of many engineering settings.

Engineering researchers can lead the innovation in mitigating rare event consequences via AI. For example, adopting, generalizing, and/or scaling the methods used to realize self-driving vehicles in a cost-effective manner captures, in large part, needed research. Engineers have always relied on modeling and simulation when empirical data is scarce. The difference in the AI context is that physics-based models must be developed alongside data-driven models to leverage the strengths of both.

There are novel and emerging technical ideas and approaches with strong potential for progress. These include deep probabilistic accelerated evaluation, model-assisted deep learning from partial observations, memory-augmented neural networks for one-shot or few-shot learning, and adaptive importance sampling for risk-based frameworks, to name a few. Innovative techniques that make efficient use of simulations in combination with expensive experimental and sparse field data are required to make effective progress on this grand challenge. In addition, the role of human behavior in AI engineered systems makes this grand challenge even more daunting, as discussed above.

Educators and engineers must develop approaches that enhance model robustness, interpretability, reliability, and causality. Additionally, researchers must incorporate constraints management and domain knowledge, leveraging cross-domain, partially incomplete, and possibly unstructured datasets. Finally, where possible, AI-assisted methods must be validated in suitable test bed environments with representative rare events and transferable findings to actual environments or suitable digital twins. These collective efforts can benefit from intelligent sampling and data-informed designs of experiments to reduce overall cost of the test and development process.

The research landscape is challenging, but it offers significant rewards. Today's AI methods are not suited for significant low-frequency or unwitnessed phenomena that engineering systems encounter. Future AI-enabled technologies that enhance awareness, certainty, response effectiveness, and reliability will lead to improved engineering systems. Stakeholders, including regulators, insurance companies, policymakers, city planners, and relevant industries, will be able to invest more effectively. Furthermore, engineers will be able to push the boundaries of design space with better understanding and mitigation of uncertainty in increasingly challenging application areas.

Incorporate Ethics in all Facets of *AI Engineering*

AI Engineering has the potential to enable rapid change at scale, but AI applications often raise new regulatory and ethical questions. Engineers must navigate these issues while ensuring compliance with all relevant laws and regulations. To this end, engineering efforts are needed to define universally accepted standards and best practices for developing, testing, and deploying AI systems.

It is important that ethical reflection, carried out by engineers and other stakeholders, keeps pace with this dynamic field. For these reasons, developing an insightful and practical tool to assist in ethical reflection upon AI-enabled engineered systems is essential. The ethical matrix—a tool developed to address related concerns in biotechnology—offers an excellent starting point for the development of such a tool.

The ethical matrix is a practical, pluralistic tool, drawing from traditions that focus on promoting wellbeing, autonomy, and justice as fairness. It encourages users to examine matters systematically, considering the point of view of each affected group.

Thus, ethical reflection on a new technology may be accomplished using the following matrix:

Stakeholder Group	Wellbeing	Autonomy	Justice
S_1	[Is S1 harmed? Benefitted?]	[Is S1's capacity for self-governance respected? Undermined?]	[Does S1 receive an unfair share of benefits? Burdens?]
S_2	[Is S2 harmed? Benefitted?]	[Is S2's capacity for self-governance respected? Undermined?]	[Does S2 receive an unfair share of benefits? Burdens?]

The matrix can make ethical reflection more efficient and comprehensive. It does not render a final verdict—unfortunately, there is no algorithm for this—but it helps bring into sharp relief the key issues and considerations for everyone involved, from researchers to designers, executives, policymakers, and the public.

While the ethical matrix is a useful tool across many domains, it can and should be tailored to a domain to better facilitate ethical reflection within that domain. This raises the question of what an ethical matrix for AI engineered systems should look like; such a matrix is both something to draw upon and continuously improve. It is a tool for use in an ongoing area of *AI Engineering* research.

AI engineered systems raise ethical concerns with respect to the following:

- **Intelligibility.** AI engineered systems will be capable of making decisions for themselves and acting on those decisions. Individuals interacting with such systems have a legitimate interest in understanding the decisions and actions of such systems.
- **Accountability.** Matters of responsibility are complicated when AI systems are involved. It is important to engineer avenues for human agents to appeal the decisions and actions of automated systems.
- **Privacy.** AI engineered systems will rely upon and possibly collect large amounts of data, raising concerns about the use of that data and concerns about privacy.
- **Bias.** Systems that rely on large bodies of data are vulnerable to absorbing biases encoded in the data.

These considerations inspire the elements for an ethical matrix for AI engineered systems:

Stakeholder Group	Intelligibility	Accountability	Privacy	Bias
S_1	...	...	...	...

However, this matrix should not be used without the original ethical matrix for two reasons. First, intelligibility, accountability, privacy, and bias only matter because each addresses the higher-level tenets of wellbeing, autonomy, and/or justice. In responding to concerns pertaining to these issues, it is vital to keep in view the fundamental values that animate those concerns. Putative concerns at this level must always be traced back to something more fundamental, such as autonomy. Second, while the matrix for AI engineered systems will assist decision makers in more quickly bringing important issues to the fore, it may not be comprehensive. To address this issue, it is helpful to also consult the original matrix, helping decision makers check their biases and also identify new columns for the matrix for AI engineered systems.

Ethical considerations for *AI Engineering* should be organically integrated into all research and education programs. This may be accomplished by inviting ethicists into these projects and activities. Equally important, specific and concrete ethical considerations should be intentionally and explicitly incorporated into all aspects and phases of the *AI Engineering* process: conceptualization, research, development, design, manufacturing, operations, and end-of-life of the system.

#3 National Initiatives for *AI Engineering*

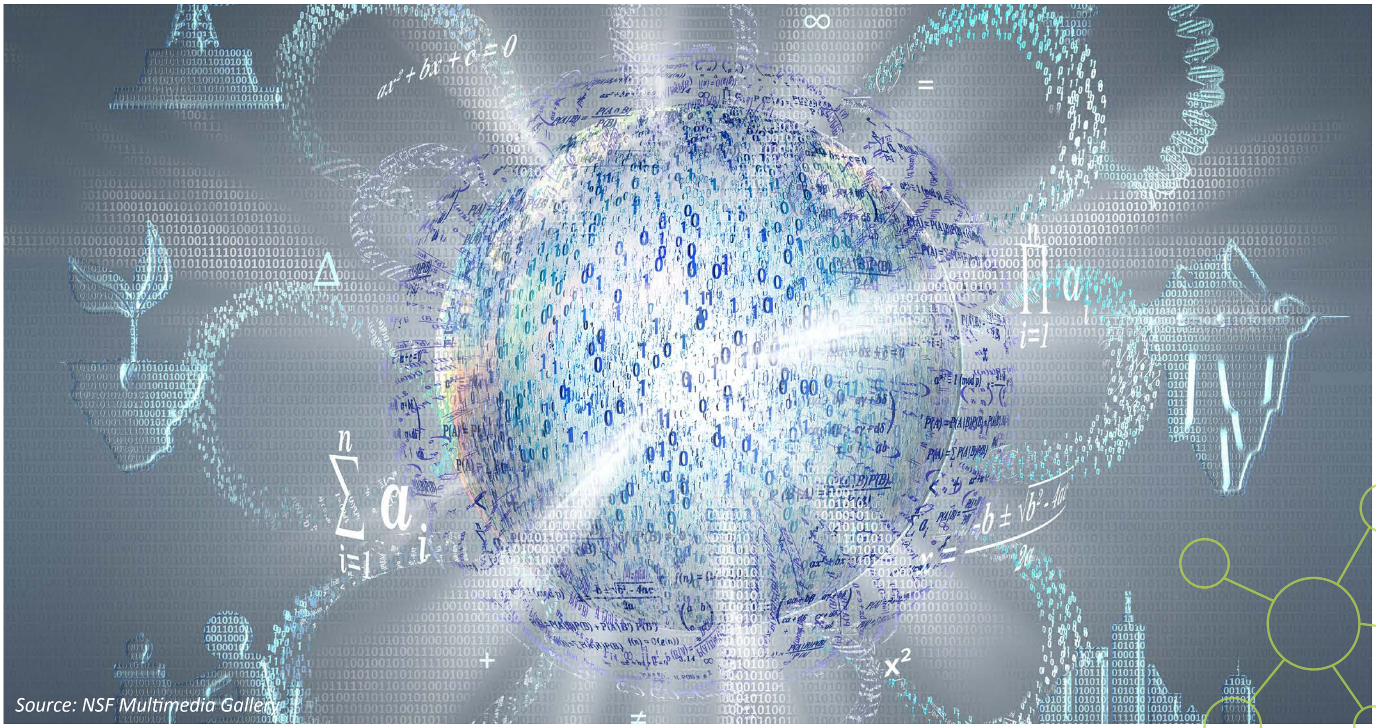

Source: NSF Multimedia Gallery

Realization of the full potential of *AI Engineering* will require large, coordinated, and sustained effort. Current efforts are well below the scale of resources needed – financial, human, organizational, and technological. Several grand challenges and opportunities explored below will require coordinated and collaborative action by leaders in academia, industry, and government at a national scale, and proactive leadership to break down silos, barriers, and boundaries will be critical for success.

Enable Collection, Curation, and Sharing of Datasets to Advance *AI Engineering*

Development of datasets is traditionally performed at the level of individual research groups. Datasets produced in this manner indeed add up to an appreciable volume. The downside is that each individual dataset is affected by changes in funding and an idiosyncratic approach by any given research group to the scope and execution of the research task. *De novo* created datasets could avoid these problems if the task is shifted to the level of **coordinated research communities**. Given the complexity and scale of the data in engineering research, decisions about data acquisition should be informed by AI systems to ensure that datasets with the highest information content are produced at minimal costs. Ideally, community-level initiatives should:

- Develop decision-making algorithms and systems (AI in the classical sense) that determine what data should be collected;
- Establish required infrastructure, computational or experimental, at the community level; and
- Create, validate, and fail-proof the relevant datasets.

This structure allows for the introduction and maintenance of a global standard for data quality, including provenance and uncertainty quantification. The datasets will be expanded systematically without creating stranded or abandoned datasets. Importantly, the data generation and consumption pattern will shift from “one group – one dataset” to a more efficient and sustainable “multiple groups – one dataset.”

Self-motivated groups of professionals sharing common goals achieved great success in the development of open-source computational tools. The underlying drivers can and should be harnessed to extend open-source community efforts into the development of datasets as well. At a fundamental level, all stakeholders will see positive outcomes. The academic community will enjoy higher efficiency in development and deployment of robust data-centric scientific processes; industry will benefit from the ability to increase mindshare and technology adoption; and government will have the best approach to public resource allocation. On the technical side, existing AI technologies offer multiple options in driving the data acquisition at a community scale. For example, automation of the computational workflows coupled with sequential decision-making, such as active learning or Bayesian optimization, can be used to create large-scale *de novo* datasets and verify existing datasets using high-performance computational facilities. At the same time, rapid evolution of robotic experimental equipment leads to the vision of **community-wide user facilities that generate empirical data**.

Practical implementation of this approach begins with the community organization into the interest groups; for example, one focused on the computational data, and the other focused on the empirical data. **The primary goal of these groups will be to develop roadmaps for automated AI-driven community research facilities dedicated to the acquisition, reproduction, and prioritization of shared data**. Similar efforts are underway in scientific research, such as Open-Source Science, which can serve as a model for the engineering research community.

Ensure Equitable Access to Computational Resources for *AI Engineering*

AI Engineering research needs a different infrastructure than traditional high-performance scientific computing, most notably heterogeneous processing using CPUs, GPUs, custom AI processors, and programmable logic devices (FPGAs). Yet, it is economically infeasible to supply that infrastructure separately for every engineering project, researcher, or student. To advance *AI Engineering* research among existing faculty, new faculty hires, and students, new resources for **research infrastructure and support** are needed at levels commensurate to that currently dedicated to both high-performance computing (HPC) and desktop computing combined. Meeting these needs requires major investments in computing resources.

While the National AI Research Resource (NAIRR) aims to produce a national resource for all AI research, the anticipated scope and pervasiveness of *AI Engineering* research will necessitate resources at the level of inter-university consortia, universities, schools of engineering, departments, research groups, industry, and individuals. Beyond computing resources, storage infrastructure and staff support (including research staff experts in methods and data engineers) must be provisioned across these levels.

Because these resources will transform and empower research and education, academic engineering research leaders should focus their efforts to secure them. This transformation of the engineering landscape can be enabled by funding programs at national and state levels to supercharge the nation’s competitiveness. Engineering research leaders in labs and companies in non-high-tech engineering sectors such as manufacturing, chemicals, transportation, and construction will need to transform their computing infrastructure. Beyond work in their respective domains and organizations, academia, industry, government, nonprofits, and civil society must collaborate across domains and sectors to adequately address this grand challenge.

Beyond structures that currently provide HPC resources effectively, the notion of a “core lab” for *AI Engineering* should be explored, like core labs in microscopy, sequencing, and animal care in biology and chemistry.

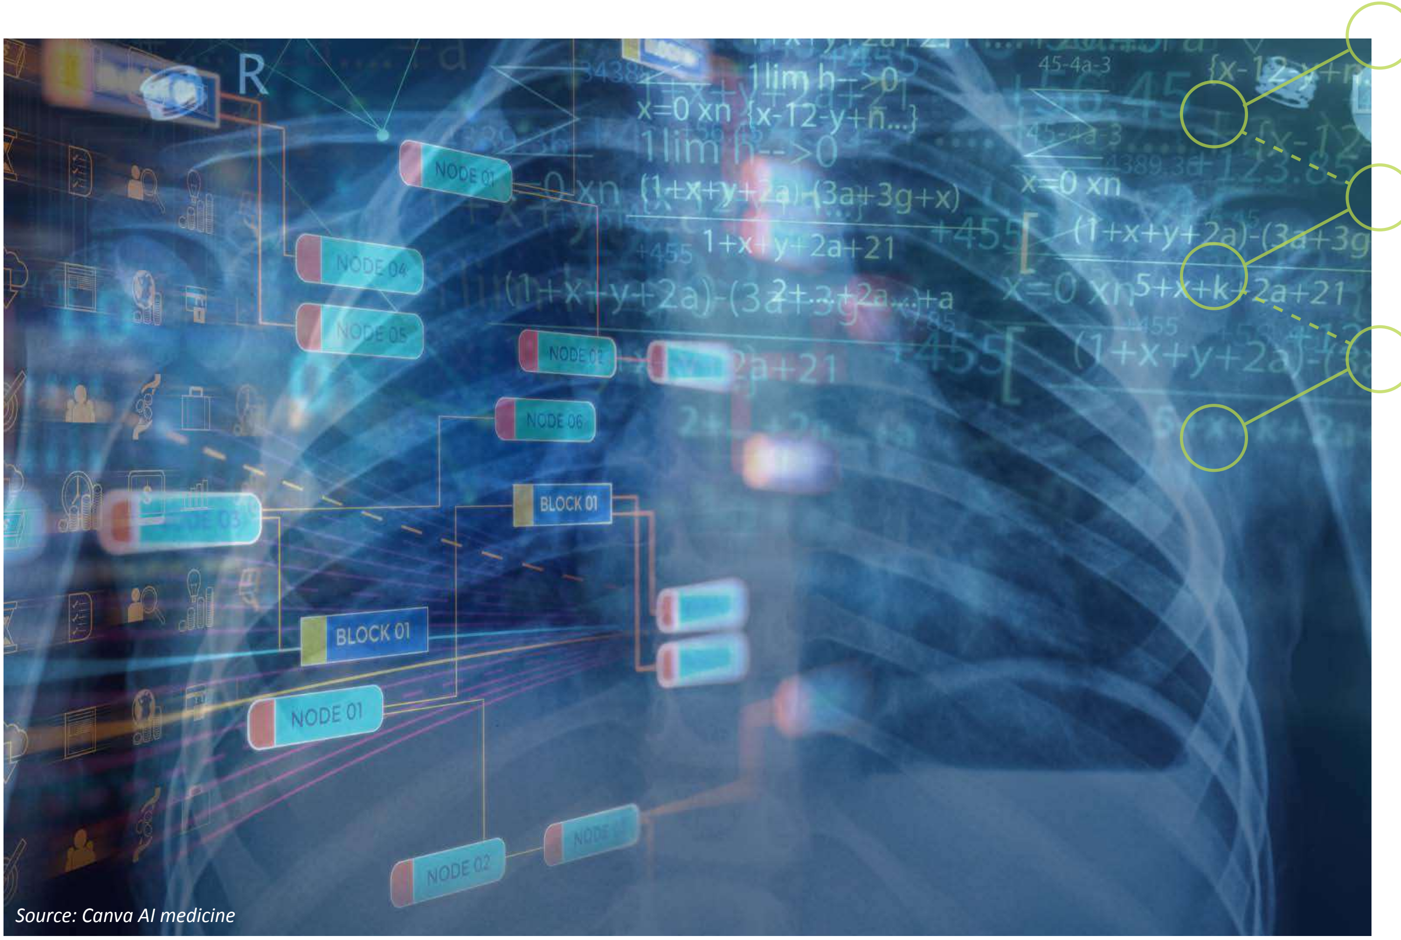

Source: Canva AI medicine

Develop Engineering Domain-Specific Foundation Models

Foundation models have emerged as one of the most promising means by which AI will have substantive impact across our social, cultural, and economic lives. These models are trained on large datasets and generate meaningful representations of the system that are compact in capturing key relationships that can be exploited for multiple downstream tasks. Society has witnessed the successful creation of first-generation foundation models in two specific applications: image and text data domains. Building on novel computational architectures such as transformers and convolutional layers, along with algorithmic innovations including self-supervised and transfer learning, foundation models are able to exploit vast quantities of unlabeled data to create large-scale, general-purpose representations capable of being fine-tuned to address a broad range of specific problems, such as generation, classification, segmentation for images (e.g., DALL-E, NOOR) and translation, summation or information extraction for natural language processing-related tasks (e.g., BERT, GPT series). In seeking to develop analogous tools to address engineering problems, substantial fundamental research will be required to create the next generation of architectures and algorithms capable of capturing the relevant patterns in the far more heterogeneous types of data encountered in problem areas, ranging from enhancing our energy infrastructure to sustainable manufacture of materials or the design and fabrication of sub-nanometer integrated circuits.

Achieving these goals will require building on the two specific algorithmic inventions that enabled the development of computationally efficient models for images and text data. **Convolution layers** in deep neural networks that capture the spatial relationships at various scales present in images have resulted in a significant improvement in the ability to classify images, and modern architectures containing convolution layers have evolved to allow for efficient computation during training on large datasets. For text data, the ability to capture sequential context has long been built on recurrent memory units, with the **attention architecture** providing the flexibility required to capture context. In both application spaces, self-supervised learning was a game changer in the efficiency of training very large networks to generate powerful foundation models.

The success of both convolutional layers and transformers rests heavily on the regular structure of their associated data types: lexicographic or temporal ordering for 1-D series or regular (typically rectilinear) grids for pixels in an image. To build foundation models that generalize across engineering applications, significant research is required for more general constructs to capture patterns exhibited in the wide range of data encountered in these fields; in addition to text, image, and video, we are provided sensor data, simulation data, and system (physical and informational) structure data. Graph-based methods, which are suited to modeling direct relations between pairs of entities, will certainly provide an initial approach that will likely prove useful in extending foundation models. It is likely that higher-order relations, as captured in hypergraphs or the more structured notion of simplicial complexes, will also have key roles to play in handling the multi-modality of systems as well as the multi-scale nature of the relevant relationships.

Foundation models thus far have been trained on vast, labeled datasets, which are not typically available for engineering applications. Indeed, the **lack of organized and available data at scale** is pervasive in engineering research; the absence of a training data pool collected over decades through content publication on accessible platforms generates another challenge.

To deal with these difficulties requires significant innovation in several areas. First, legacy data must be more effectively used. Developing methods and tools for rigorously addressing issues, including quantity, quality, labels, and completeness, will be required to make the best possible use of the existing data. Second, data collation from multiple sources can provide datasets of sufficient size for training foundation models. Concepts like federated learning are a possible approach to balance information sharing and privacy, but the application to similar yet different generator systems proves to be another challenge. Developing algorithms to preserve data privacy upon sharing is another path to compile datasets with sufficient variety and volume that allows training of foundation models.

New training paradigms and computational architectures must be developed. Purely data-driven models must be modified to efficiently incorporate fundamental domain knowledge available in various forms. Physics-based laws (equations of Newton, Maxwell, Navier-Stokes, diffusion, radiative transfer, etc.), associated constitutive relations (Laws of Ohm, Fick, and Hooke, etc.), conservation principles (mass, energy, linear momentum, angular momentum), and symmetries (invariances to shift, rotation, scale, as well as more general Lie-based and gauge transformations) provide critical constraints on engineered systems. The pursuit of incorporating these principles into AI models has only recently begun. Finally, fundamentally new architecture with far fewer parameters will be needed.

For foundation models to truly be useful across engineering fields, the **issue of uncertainty** must be addressed. Uncertainty is ubiquitous within the context of built systems. Many modeling structures have been developed to address uncertainty, including probabilistic approaches and worst-case analysis. It is easy to model a few discrete parameters in classical settings, while it is much more challenging to describe high- or infinite-dimensional (i.e., functional) unknowns. How to do all of this in the context of an AI foundation model is uncharted territory but very much needed.

A final, critical aspect in the quest to generate foundation models concerns the **trustworthiness of predictions or modeled outputs**. Since scientific systems in engineering are not free of constraints, it is important to focus on the capacity and tools to interpret why objectives may not be met due to constraints and/or lack of reliable training data. The ability to provide guidance to the user as to when and why predictions and generations are trustworthy is a key requirement for meaningful foundation models.

We believe this grand challenge should be addressed at the national level to ensure access to such foundation models to the full range of the research community rather than to a few institutions or companies. Strategic collaboration involving federal agency leaders, academic research domain leaders, and industry research lab leaders will be crucial to develop generally available, up-to-date, secure, reliable, and trustworthy foundation models.

Establish Dedicated Research Institutes for *AI Engineering*

NSF launched the AI Institutes program in 2020, under the auspices of which 25 have been created, encompassing a wide range of fields of scientific research. These AI Institutes are expected to lay the foundations for trustworthy AI and advance our understanding of AI theory and applications.

Alongside the AI Institutes, there is a major need and unrealized opportunity to create a set of ***AI Engineering* Institutes**. These institutes would combine the best elements of NSF AI Institutes and NSF Engineering Research Centers. They will foster creation of new engineering knowledge and methodologies in both directions: engineering for AI and AI for engineering. These institutes must include collaborative representation from government, academia, and industry. For illustrative purposes, below are three examples of potential *AI Engineering* Institutes.

Institute for *AI Engineering* of Autonomous Systems. Recent advancements in AI have led to significant developments in the autonomous operation of vehicles, drones, and robotics. This evolution presents unique opportunities and challenges, particularly in ensuring the safety, reliability, and efficiency of these systems. We propose investing in research and development that focuses on the **optimization of autonomous system design**, enhancing their decision-making capabilities, and ensuring their responsible integration into society. Such investment will pave the way for safer, more efficient, trustworthy, and innovative autonomous technologies that will benefit humanity.

Institute for AI Brain Engineering. Leveraging AI to decode complex neural processes can revolutionize our understanding of the human brain. This includes developing computational models to simulate brain functions and applying AI to analyze neurological data. Such advancements can lead to breakthroughs in treating brain disorders and enhancing cognitive abilities. Investment in this field promises significant contributions to neuroscience and psychology, aligning technological progress with a deeper comprehension of human cognition and behavior. At the same time, understanding how the brain senses, processes, and prioritizes sensory information to guide behavior will also lead to new bio-inspired computing architectures.

Institutes for *AI Engineering* Science. Several national and international reports have highlighted the opportunity for AI to revolutionize scientific discovery. An early example of such a breakthrough was using AI techniques to solve the protein folding problem. More recent examples include discovery of new antibiotics and new battery materials. Automating the processes that enable the discovery of new materials and characterization of new phenomena can lead to a transformational AI-centric approach to discovery, modeling and characterizing physical phenomena.

Some modalities for connecting AI to science are, in principle, applicable to engineering science fields. NSF has already launched a program of AI Institutes for materials science and engineering. **New AI Institutes that can address other fundamental fields of engineering science** (e.g., thermal sciences, fluid dynamics, chemical engineering, polymer sciences, quantum engineering, environmental science and engineering, structural engineering, power systems, etc.) could enable dramatic advances in these engineering fields. In addition, fundamental problems from engineering design, manufacturing, and operations have influenced fields of engineering science, e.g., control theory, signal processing, communications, etc., and therefore are potential topics for this class of *AI Engineering* Institutes.

Industrial Data Foundry for Accelerated AI Development and Validation

The Industrial AI Data Foundry will serve as a hub to host different datasets from a broad range of industries, including automotive, aerospace, health care, semiconductors, energy, transportation, manufacturing, and industrial automation. The Industrial AI Data Foundry will comprise the critical mass investments that are expected to form connections to the wider AI community and research and innovation ecosystem. The goal of the Industrial AI Data Foundry is to make data available to researchers, analysts, developers, and other interested parties (including industry and government agencies) for the purpose of facilitating research, innovation, and discovery.

Create New Education and Training Programs for *AI Engineering*

For *AI Engineering* to thrive, developing new education and training programs in this emerging field is vital at all levels: undergraduate, graduate, and professional education. Implementing AI in various domains requires certain levels of skills and expertise that are lacking at scale in the United States and worldwide. Significant work is needed to determine how to expand educational efforts to meet the growing demand for people who can not only create AI but also manage it, especially in safety-critical settings. Current education efforts do not address the bold initiatives required to advance engineering education at all levels with a view to *AI Engineering* education for the coming decade. To make a major impact on AI for engineering, we need an engineering-focused consortium across government, universities, industry, civil society, and nonprofits with a clear set of goals.

Engineering schools and colleges are ideally positioned to capitalize on the opportunity to collaborate, along with industry and government. These programs can build on existing strengths in data science, machine learning, and related topics; sharing materials and resources across institutions should also be encouraged. Elements of these courses and programs can be useful to current faculty and research staff. The visioning event was focused on engineering research and did not have education and workforce development as core topics for discussion. As such, the discussion that follows is far from sufficient. The urgency of this topic requires an immediate focused initiative on *AI Engineering* education.

Undergraduate Education

At a minimum, undergraduate students in all engineering fields will require a basic **course in AI fundamentals**. Although adding a new requirement in undergraduate degree programs is extraordinarily difficult, engineering educators, leaders, and ABET should consider infusing AI into their curriculum. This undergraduate exposure could come via a common course, by specialized courses by major, or some combination thereof. This approach is analogous to how a requirement in computing or statistics is currently met. The proposed AI requirement will add burdens on teaching resources, which can be met creatively by using public online resources.

It is clear that traditional undergraduate engineering curriculum will need to adapt with considerable urgency as *AI Engineering* develops in the coming years. As generative AI-based tools proliferate across various knowledge domains, engineering educators need to confront the core issue of how to teach engineering fundamentals. Even the role, structure, and duration of the traditional four-year undergraduate degree requires thoughtful and open discussion. A dedicated workshop on this topic may be required.

AI presents many great experiential learning opportunities that could be coordinated between industry and academia.

- Companies could provide summer internships employing methods taught the previous semester.
- Industry could provide sanitized data and associated problems for AI fundamentals classes or more advanced classes, especially capstones.

Undergraduate majors are launching in *AI Engineering*, such as that announced by the University of Pennsylvania in February 2024, but it is critical that students in all areas of engineering understand the impact of AI on their work.

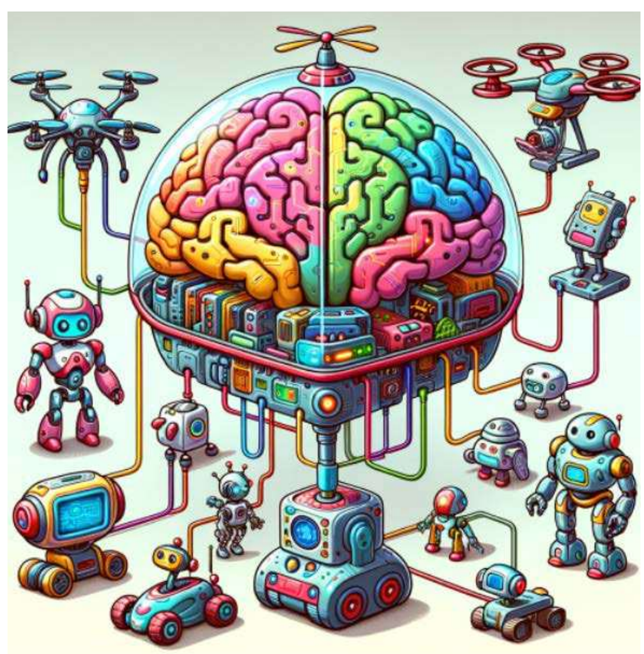

Source University of Pennsylvania, A DALL-E generated image to illustrate a robot brain, general learning algorithms for embodied AI.

Graduate Education

For graduate education, each engineering field should devise courses for both master's and doctoral students. Some could be shared across departments while others will be more specialized to the field. Graduate engineering education should develop suitably revamped and differentiated curricula in *AI Engineering*, with one path for those who plan to go into industry and another for those who plan to enter research. Both of these directions will be profoundly impacted by generative AI technologies. Graduate-level courses combining data science, machine learning, and AI with engineering science courses from signal processing, information theory, optimization methods, control theory, etc., should be developed, especially for doctoral students. Developing master's certificate and/or master's degree programs in *AI Engineering* would also support education and talent development needs.

Professional Education

A substantial opportunity exists to develop optimized short courses and certificate programs for industry professionals to rapidly learn the aspects of *AI Engineering* relevant to their work. Such courses and programs should be developed in collaboration with industry and government. Specialized short courses should be developed for engineering professionals, technology managers, and leaders in government agencies.

Conclusion

AI Engineering is an enormously compelling opportunity for the nation to maintain and enhance U.S. global leadership in engineering and technology and supercharge U.S. economic competitiveness, national security, and public health. It will enable new discoveries and inventions, accelerate design, transform manufacturing efficiency and sustainability, and enhance operations of engineered products, systems, and processes. Over the next five to 10 years, the entire landscape of engineering, technology, economy, and society will be substantially different from what might be imagined in a business-as-usual scenario.

This report has outlined major themes, grand challenges, and opportunities in the emerging landscape of *AI Engineering*. They will likely remain verdant for some years to come, but AI is a rapidly developing field, and major breakthroughs in the coming years are anticipated. With this in mind, the specific details of the grand challenge opportunities in this report may need to be reassessed on a periodic basis.

Realization of the full potential of *AI Engineering* requires convergence, coordination, and collaboration of people and organizations from academia, industry, and government. Current efforts fall far short of what is necessary; challenges are formidable, and global competition is fierce. It is essential to mobilize large-scale financial, technological, human, and organizational resources now to meet rapidly expanding needs. This will not happen without strong, proactive, coordinated, and collaborative action by leaders working across sectors to enable *AI Engineering*.

Appendix A: Visioning Event Participants

Ayomide Afolabi, Kennesaw State University Saurabh Amin, Massachusetts Institute of Technology

Ella Atkins, Virginia Polytechnic Institute

Michael Branicky, University of Kansas

Birgit Braun, Dow Chemical Co.

Clinton Castro, University of Wisconsin-Madison

Jill Crisman, Underwriters Laboratories Research Institutes

Missy Cummings, George Mason University

Duane Detwiler, Honda Research Institute USA

Charles Farrar, Los Alamos National Laboratory

Brian Giera, Lawrence Livermore National Laboratory

Jennie Hwang, H-Technologies Group

Mahesh Iyer, Intel Corporation

H.V. Jagadish, University of Michigan

Henry Jerez, Google

Yunji (Rosie) Kim, Georgia Institute of Technology

Vijay Kumar, University of Pennsylvania

Jay Lee, University of Maryland College Park

Manuel Martinez-Ramon, The University of New Mexico

Benji Maruyama, Air Force Research Laboratory

Eric Miller, Tufts University

Sanghamitra Neogi, University of Colorado Boulder

Onyema Osuagwu, Morgan State University

Rohit Pandeka, Microsoft

Kayla Partee, Clark Atlanta University

Ying Qian, University of Georgia

William Sanders, Carnegie Mellon University

Pascal Van Hentenryck, Georgia Institute of Technology

Sophia Velastegui, Aptiv

Yael Vodovotz, The Ohio State University

Dmitry Zubarev, IBM Research

National Science Foundation Participants

Eyad Abed

Louise Howe

Crystal Leach

Reha Uzsoy

Contractual Partners

Susan Lang, The Ohio State University

David Glazer, Glazer Design

ERVA Team and Consultants

Josh Aebischer, ERVA

Mike Brizek, UIDP

Dorota Grejner-Brzezinska, The Ohio State University

Pramod Khargonekar, University of California, Irvine

Sandy Mau, UIDP

Rebecca Silveston, ERVA

Appendix B: Event Guest Speaker Presentation Summaries

During the visioning event, ERVA invited two distinguished voices in AI to present their perspectives on the research landscape and the potential for *AI Engineering*.

Both speakers concluded their talks with the observation that AI must be developed ethically and responsibly. Consumer-facing risks, including disinformation and deep fakes, surveillance and privacy concerns, bias and discriminatory algorithms, social and behavioral engineering, and market manipulation, are among acknowledged hazards. Researchers must also consider security risks that AI could exacerbate, including those tied to nuclear, chemical, and cyber warfare. To create trustworthy and reliable AI, researchers must consider how to compare the behavior of different models and how to assess domain knowledge, emergent behaviors, and knowledge synthesis, among other factors.

Multimodal Biomedical AI

Eric Topol, Founder and Director, Scripps Research Translational Institute; Executive Vice President, Scripps Research; Professor, Molecular Medicine, Scripps Research; Gary and Mary West Endowed Chair of Innovative Medicine, Scripps Research

Topol explained how AI has transformed and continues to transform the medical profession. He focused primarily on the area of medical errors. He noted that while most people will experience at least one diagnostic error in their lifetime, over 800,000 people die or suffer serious consequences annually from these misdiagnoses.

AI can help make medicine more precise and accurate. One example, which Topol termed "machine eyes," is unimodal, supervised AI, and has rapidly advanced in the last six years. Whether accessing retina scans, x-rays, or cardiograms, AI, which can read these images beyond the capability of humans to predict a variety of medical conditions, from blood pressure and glucose levels to diabetes, Alzheimer's, Parkinson's, and gall bladder disease. This type of AI has also been used to streamline workflows in routine diagnostic procedures such as mammograms and colon cancer screening. It has also been used to accelerate the diagnosis of critically ill infants, moving from tissue sample to diagnosis in 13.5 hours—enough time to prevent severe organ damage or other debilitating conditions.

Supervised AI, though, has limitations, as humans must code the input data needed to train the model. The current challenge for medicine is to capitalize on the potential of transformer model architectures, as well as self-supervised and unsupervised learning, in the creation of multimodal AI. This AI uses data from a variety of sources, including both those commonly available to clinicians and those less accessible to most. Advances in sensors, coupled with transformer AI models, can potentially provide "pre-womb to tomb" health care, as we may now capture each individual's physical data and integrate it with health records, environment, and social factors to perform many data-driven tasks. Existing virtual health assistants could add a preventative or holistic element currently missing in the health care landscape. Multimodal AI could enable the concept of the "hospital at home" to become a reality with the ability to anticipate patient deterioration and shorten or avoid unnecessary hospital visits altogether. AI, combined with smartphone apps that enable what Topol called "medical selfies," can provide people with more agency in their health care.

A Case for Sustainability

Rick Stevens, Professor of Computer Science, University of Chicago; Associate Laboratory Director, Computing, Environment and Life Sciences, Argonne National Laboratory

Stevens explained how a 2022 series of developmental workshops organized by the Department of Energy national laboratories under the guidance of the Office of Science (SC) and the National Nuclear Security Administration (NNSA) considered six cross-cutting themes focused on AI. A report synthesizing findings from the series points to the critical importance of harnessing AI to address national imperatives in advanced materials properties and inverse design, robotics for automated discovery, surrogates for high-performance computing, software engineering and programming, predictive components for complex engineered systems, and foundations for scientific knowledge. He noted that advancing in these areas would be possible only through collaboration between government agencies, industry, academia, and international partners.

He then focused on the need for development of foundation ecosystems that could be used for many different scientific tasks, noting that such ecosystems are not likely to be developed by anyone other than scientists, given the inherent limitations on short-term financial gain. Stevens noted that tasks that could be handled by such a model—summarizing and distilling knowledge from a dataset—could accomplish in a weekend what it might take a team of humans several years or more to complete. Such models could also synthesize information and generate small computer programs to accomplish specific tasks. They could generate plans, solve logic problems, and write experimental programs for robots. More research and development is needed, however, before foundation models might approach artificial general intelligence capable of generating useful, verifiable hypotheses and theories for testing.

Further Reading

Agrawal, A. K., Gans, J. S., & Goldfarb, A. (2023, October 9). The Turing Transformation: Artificial Intelligence, intelligence augmentation, and skill premiums. National Bureau of Economic Research. nber.org/papers/w31767

Artificial Intelligence in science. OECD iLibrary. (2023). oecd-ilibrary.org/science-and-technology/artificial-intelligence-in-science_a8d820bd-en

Birhane, A., Kasirzadeh, A., Leslie, D., & Wachter, S. (2023, April 26). Science in the age of large language models. Nature News. nature.com/articles/s42254-023-00581-4

Carter, J., Feddema, J., Kothe, D., Neely, R., Pruet, J., Stevens, R., Balaprakash, P., Beckman, P., Foster, I., Iskra, K., Ramanathan, A., Taylor, V., Thakur, R., Agarwal, D., Crivelli, S., de Jong, B., Rouson, D., Sohn, M., Wetter, M., … Dietrich, E. M. (2023, May 1). Advanced Research Directions on AI for science, energy, and security: Report on Summer 2022 workshops. (Technical Report) | OSTI.GOV. osti.gov/biblio/1986455

Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023, June 14). The economic potential of Generative AI: The Next Productivity Frontier. McKinsey & Company. mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#introduction

De Peuter, S., Oulasvirta, A., & Kaski, S. (2023, March 3). Toward AI assistants that let designers design. onlinelibrary.wiley.com/doi/full/10.1002/aaai.12077

Engineering Research Visioning Alliance, Engineering the Future of Distributed Manufacturing: A Visioning Report, 2023. SSRN. dx.doi.org/10.2139/ssrn.4512398.

Engineering Research Visioning Alliance, Engineering Materials for a Sustainable Future: A Visioning Report, 2024. SSRN. dx.doi.org/10.2139/4663931.

Gao, J., & Wang, D. (2023, April 17). Quantifying the benefit of Artificial Intelligence for Scientific Research. arXiv.org. https://arxiv.org/abs/2304.10578

Gu, B. (2023, November 14). Optimizing Electronics Design with AI co-pilots. EE Times. eetimes.com/optimizing-electronics-design-with-ai-co-pilots/

Kanarik, K. J., Osowiecki, W. T., Lu, Y. (Joe), Talukder, D., Roschewsky, N., Park, S. N., Kamon, M., Fried, D. M., & Gottscho, R. A. (2023, March 8). Human–machine collaboration for Improving Semiconductor Process Development. Nature News. nature.com/articles/s41586-023-05773-7

Kim, J. (2023, December 6). Google Deepmind's new AI tool helped create more than 700 new materials. MIT Technology Review. technologyreview.com/2023/11/29/1084061/deepmind-ai-tool-for-new-materials-discovery/

Kossiakoff, A., Biemer, S. M., Seymour, S. J., & Flanigan, D. A. (2020, July 8). Systems engineering principles and practice, 3rd Edition. Wiley.com. wiley.com/en-us/Systems+Engineering+Principles+and+Practice,+3rd+Edition-p-9781119516668

Lara, F. and Deckers, J., eds. (2024). Ethics of Artificial Intelligence. Vol. 41. SpringerLink. link.springer.com/book/10.1007/978-3-031-48135-2

Madni, A. M. (2018). Transdisciplinary Systems Engineering. SpringerLink. https://link.springer.com/book/10.1007/978-3-319-62184-5

Miller, C. (2024, January 11). Global chip war for strategic semiconductors. Nature News. nature.com/articles/s44287-023-00005-6

National Academies of Sciences. (2023, September 7). Test and evaluation challenges in artificial intelligence-enabled systems for the Department of the Air Force. The National Academies Press. nap.nationalacademies.org/catalog/27092/test-and-evaluation-challenges-in-artificial-intelligence-enabled-systems-for-the-department-of-the-air-force

National Artificial Intelligence Research Resource Task Force. (2023). Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem - An Implementation Plan for a National Artificial Intelligence Research Resource. ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf

National Network for Critical Technology Assessment. (2023). Report. Securing America's Future: A Framework for Critical Technology Assessment. nncta.org/opportunities/report.html

NSPE Code of Ethics. National Society of Professional Engineers. (2019). nspe.org/resources/ethics/code-ethics

National Science and Technology Council. The National Artificial Intelligence R&D Strategic Plan 2023 Update, Office of Science and Technology Policy, Executive Office of the President, USA, 2023.

National Science Foundation. Democratizing the future of AI R&D: NSF to launch national AI research resource pilot, NSF News, 2024. new.nsf.gov/news/democratizing-future-ai-rd-nsf-launch-national-ai

Nti, I. K., Adekoya, A. F., Weyori, B. A., & Nyarko-Boateng, O. (2021, April 15). Applications of artificial intelligence in engineering and manufacturing: A systematic review - Journal of Intelligent Manufacturing. SpringerLink. link.springer.com/article/10.1007/s10845-021-01771-6

O'Neil, C., & Gunn, H. (2020, September 17). Near-term Artificial Intelligence and the ethical matrix. OUP Academic. academic.oup.com/book/33540/chapter-abstract/287905687?redirectedFrom=fulltext

Perry, T. S. (2023, November 15). How will AI affect the semiconductor industry? IEEE Spectrum. spectrum.ieee.org/how-will-ai-change-semiconductors

Sabherwal, R., & Grover, V. (2024). The societal impacts of Generative Artificial Intelligence: A balanced perspective. AIS Electronic Library (AISeL). aisel.aisnet.org/jais/vol25/iss1/14/

Sakao, T., Funk, P., Matschewsky, J., Bengtsson, M., and Ahmed, M. U. AI-LCE: Adaptive and Intelligent Life Cycle Engineering by applying digitalization and AI methods–An emerging paradigm shift in Life Cycle Engineering. Procedia CIRP, Volume 98, 2021, pages 571-576, ISSN 2212-8271, doi.org/10.1016/j.procir.2021.01.153.

Schmidt, E., Work, B., Catz, S., Chien, S., Darby, C., Ford, K., ... & Matheny, J. Final Report of the National Security Commission on Artificial Intelligence, 2021.

Semiconductor Research Corporation. Decadal Plan for Semiconductors, 2023.

Turing 2.0: Changing the world for the better with data science and AI. The Alan Turing Institute, 2023.

The United States Government. (2023, October 30). Executive order on the safe, secure, and trustworthy development and use of artificial intelligence. The White House. whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

Wang, H., Fu, T., Du, Y., Gao, W., Huang, K., Liu, Z., Chandak, P., Liu, S., Van Katwyk, P., Deac, A., Anandkumar, A., Bergen, K., Gomes, C. P., Ho, S., Kohli, P., Lasenby, J., Leskovec, J., Liu, T.-Y., Manrai, A., … Zitnik, M. (2023, August 2). Scientific discovery in the age of Artificial Intelligence. Nature News. nature.com/articles/s41586-023-06221-2

Wojcik, H. (2024, February 13). Penn Engineering Announces First Ivy League undergraduate degree in Artificial Intelligence. Penn Engineering Blog. blog.seas.upenn.edu/penn-engineering-announces-first-ivy-league-undergraduate-degree-in-artificial-intelligence/

Yin, J., Dash, S., Wang, F., & Shankar, M. (2023, November 1). Forge: Pre-training open foundation models for science: Proceedings of the international conference for high performance computing, networking, storage and analysis. ACM Conferences. dl.acm.org/doi/10.1145/3581784.3613215

NSF Engineering Research
Visioning Alliance

Our mission is to identify and develop bold and transformative new engineering research directions and to catalyze the engineering community's pursuit of innovative, high-impact research that benefits society.

ERVA IS FUNDED BY THE NATIONAL SCIENCE FOUNDATION THROUGH AWARD NUMBER 2048419

This material is based upon work supported by the National Science Foundation under Grant # 2048419. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Ervacommunity.org