



Testimony of

**James F. Kurose, Ph.D.
Assistant Director of the
National Science Foundation for
Computer and Information Science and Engineering**

Before the

Subcommittee on Information Technology

For the

**Committee on Oversight and Government Reform
U.S. House of Representatives**

March 7, 2018

Game Changers: Artificial Intelligence Part II

Good afternoon, Chairman Hurd, Ranking Member Kelly, and members of the Subcommittee. My name is Jim Kurose and I am the Assistant Director of the National Science Foundation (NSF) for Computer and Information Science and Engineering (CISE).

As you know, NSF is dedicated to advancing progress in all fields of science and engineering. NSF funds fundamental research across all science and engineering disciplines; supports education of the next generation of innovative thinkers, discoverers, and leaders; and contributes to national security and U.S. economic competitiveness. I welcome this opportunity to highlight NSF's investments in artificial intelligence (AI).

Federal investments in fundamental, long-term, transformative AI research as well as education are critical to achieving and sustaining U.S. technological leadership in this area. The recent *National Artificial Intelligence Research and Development Strategic Plan*¹, developed by a National Science and Technology Council (NSTC) Subcommittee on Machine Learning and Artificial Intelligence (MLAI), noted the important role of the Federal Government in the broader AI innovation ecosystem. The strategic plan identified scientific and technological needs in AI, established priorities for federally-funded research and development (R&D) in AI, and articulated a path forward that looked beyond near-term AI capabilities toward the long-term transformational impacts of AI on our Nation.

¹ https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf

Over the last several years, we have witnessed AI begin to have broad and deep impacts on our daily lives, ranging from precision medicine to intelligent transportation, to personalized education and learning, and beyond. Importantly, the many powerful AI innovations being led by industry today are predicated on the fundamental research outcomes generated with federal funding, including NSF funding, over the last several decades. Similarly, sustained investments now will lay the foundations for future innovations in the decades to come.

NSF is continuing to fund pioneering innovations that will help the U.S. capitalize on the full potential of AI to strengthen our economy, advance job growth, and better our society. Looking forward, NSF will continue to bring the problem-solving capabilities of the Nation's best and brightest minds to bear on the AI challenges of today and tomorrow.

A Brief History of AI

Let me begin with a brief history of AI. AI innovations date back to the early 1940s, when researchers constructed the first Boolean circuit model of the human brain. In 1950, Alan Turing published the seminal paper *Computing Machinery and Intelligence*², in which he considered the question, "Can machines think?" In that paper, Turing introduced his concept of what is now known as the "Turing test," or the test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. The Turing test called for a human evaluator to judge natural language conversations between a human and a machine designed to generate human-like responses; if the evaluator cannot reliably tell the machine from the human, the machine is said to have passed the test.

Early AI programs were developed in the 1950s, and the term "artificial intelligence" was officially adopted at a meeting at Dartmouth College in 1956 to reflect the study of machines capable of intelligent behavior. Development of knowledge-based systems in the 1970s led to an expert systems boom in the 1980s. That, however, was followed by a period known as the "AI winter" in the late 1980s and early 1990s, when there was reduced interest in the field. Beginning in 2005, the broad availability of very large data sets combined with advanced, scalable computing systems catalyzed a new era of deep learning, and since 2013, AI systems have begun to surpass human capabilities in image processing, speech recognition, and even the game of Go³.

Importantly, beginning in the 1960s, NSF supported the very first computer science departments at U.S. colleges and universities. Often growing out of mathematics departments, establishing computer science as a mainstream area of scientific and engineering pursuit helped to provide a fertile training ground for the first (and subsequent) generations of computer scientists and engineers, including AI pioneers. Today, NSF provides 83 percent of all federal support for fundamental research in computer science, including AI, conducted at our Nation's colleges and universities.

At the same time, the pursuit of the formal and theoretical foundations of AI has been a highly interdisciplinary endeavor, spanning many fields beyond computer science and engineering, to include philosophy, mathematics and control theory, economics and game theory, psychology, linguistics, and ethics. By bringing together these varied disciplines, NSF has been uniquely positioned to drive advances in the foundations of AI.

² <https://academic.oup.com/mind/article-pdf/LIX/236/433/9866119/433.pdf>

³ <https://www.nytimes.com/2017/05/23/business/google-deepmind-alphago-go-champion-defeat.html>

NSF's Sustained Investment in AI: The Foundations for Future Advances

In keeping with its mission “to promote the progress of science; to advance the national health, prosperity, and welfare; [and] to secure the national defense,” NSF has long supported fundamental research related to AI. Importantly, NSF prioritizes its AI investments through a proven, “bottom-up” philosophy: the best ideas for research come directly from the science and engineering community. As Eric Schmidt, former Chief Executive of Google and more recently Alphabet, Inc., Google’s parent company, has said, NSF is “where all interesting research gets started.”

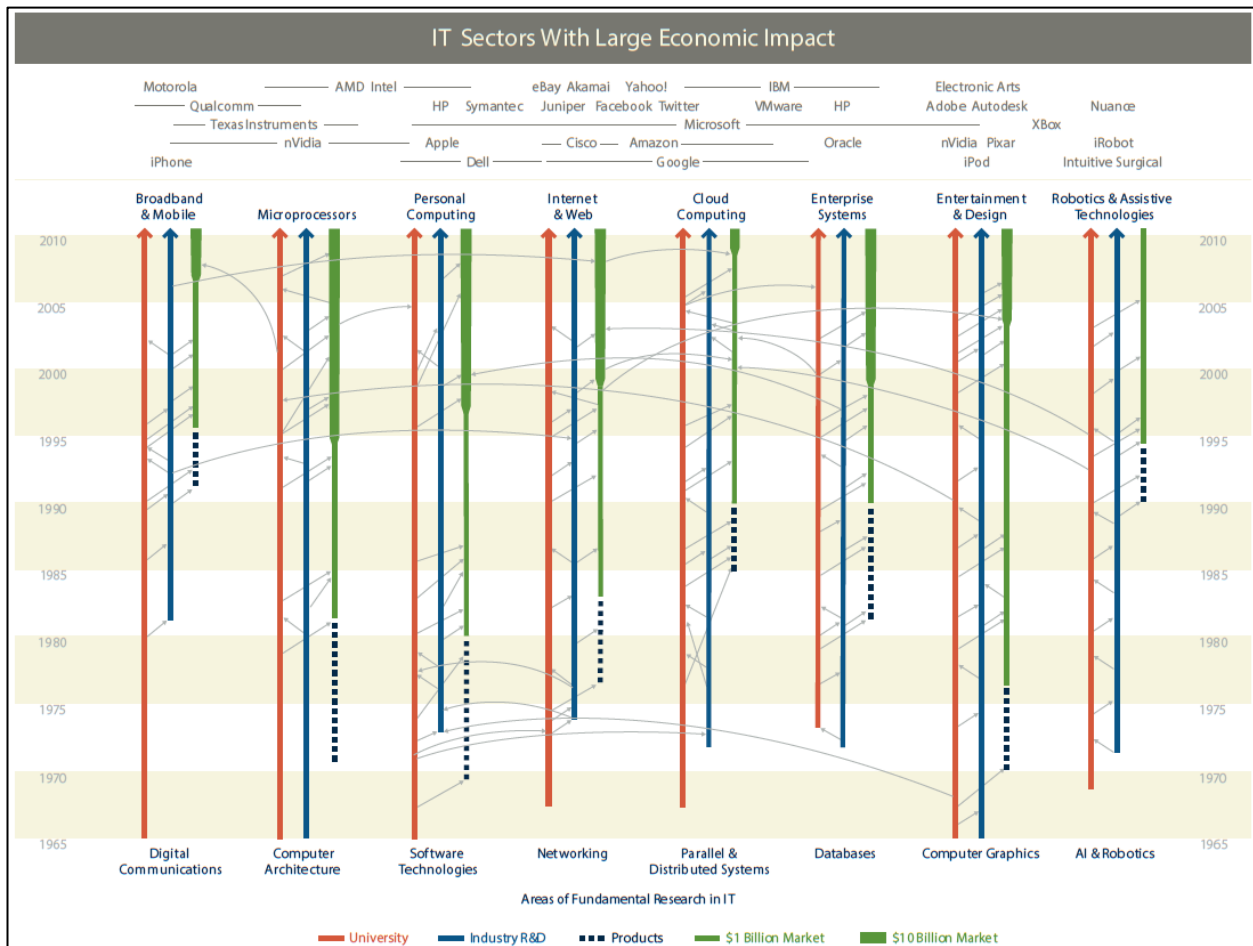


Figure 1. The innovation ecosystem for AI and robotics. From *Continuing Innovation in Information Technology Workshop Report*, National Academies of Sciences, Engineering, and Medicine, 2016.

Figure 1, a so-called “tire tracks” diagram, from the *Continuing Innovation in Information Technology Workshop Report*⁴ published by the National Academies of Sciences, Engineering, and Medicine, illustrates this point. The figure shows how fundamental research in information technology, conducted at universities with federal funding, as well as in industry, has led to the introduction of entirely new product categories that have resulted in multi-billion-dollar markets. The figure reflects the information technology research ecosystem in which concurrent advances across multiple sectors have been mutually reinforcing, stimulating, and enabling one another, and ultimately leading to a vibrant, innovative industry exemplified by top-performing firms.

⁴ <https://www.nap.edu/catalog/23393/continuing-innovation-in-information-technology-workshop-report>

Specifically, the vertical red track represents university-based (and largely federally-funded) research, and the blue track represents industry R&D (some of which is also government-funded). The dashed and solid black line indicates a period following the introduction of significant commercial products resulting from this research, and the green line represents billion-dollar-plus industries (by annual revenue) stemming from this research. The top rows list the present-day information technology market segment and representative firms whose creation was stimulated by the decades-long research represented by the red and blue vertical tracks.

Importantly, the arrows between the vertical tracks represent known instances of cross-fertilization resulting from multi-directional flows of ideas, technologies, and people across academic research, industry research, and products. In other words, they are examples of the rich interplay between academia, industry, and government that characterizes the unique American information technology innovation ecosystem that has given rise to multi-billion-dollar markets, including in the case of AI and robotics (shown on the far right of the figure).

As a concrete example of this innovation ecosystem, consider how pairing AI research with the growth of the Internet in the 1990s enabled the creation of e-commerce, a crucial driver of today's economy. NSF-funded researchers began working on what is now known as collaborative filtering, developing and refining the origins of this technique. Today, collaborative filtering fuels the recommender engines on popular websites like Netflix and Amazon – the “you might also like” suggestions that propel a significant proportion of e-commerce activity⁵.

Neural networks constitute another key innovation rooted in AI research. Although this modeling approach emerged in the late 1980s, there were not enough data available at the time for neural networks to make accurate predictions. With the rise of big data and today's data-intensive scientific methods, together with conceptual advances in how to structure and optimize the networks, neural networks have re-emerged as a useful way to improve accuracy in AI models. For example, in recent years, neural networks have helped reduce the error rate in speech recognition systems, enabling innovations such as real-time translation⁶.

Similarly, NSF's investments in reinforcement learning – an approach rooted in behavioral psychology that involves learning to associate behaviors with desired outcomes – have led to today's deep learning systems. By getting computers to learn like humans, without explicit instruction, reinforcement learning is driving progress in self-driving cars and other forms of automation where machines can hone skills through experience. Reinforcement learning was the key technology underlying AlphaGo, the program that defeated the world's best Go players⁷.

These are but a few of the many NSF-funded AI innovations that impact our lives today.

⁵ <https://www.nap.edu/catalog/23393/continuing-innovation-in-information-technology-workshop-report>

⁶ <https://www.nap.edu/catalog/23393/continuing-innovation-in-information-technology-workshop-report>

⁷ <https://www.technologyreview.com/s/603501/10-breakthrough-technologies-2017-reinforcement-learning/>

The Frontiers of AI Research

Going forward, long-term, fundamental AI R&D challenges can be classified into two categories known as “narrow AI” and “general-purpose (general) AI.” Narrow AI is focused on solving specific tasks in well-defined domains, such as speech recognition or computer vision. Examples of narrow AI advances include facial recognition and AlphaGo. Narrow AI R&D challenges include robustness, adaptability, trust and safety, taskability, interaction and collaboration, knowledge richness, and scalability. By contrast, general AI is about exhibiting flexibility and versatility in a broad range of cognitive domains, including learning, reasoning, creativity, and planning; it involves transferring what is learned or experienced in one task to another, and ultimately appreciating “intent,” “meaning,” and “understanding” in AI systems. General AI R&D challenges span cognition and metacognition, visual intelligence, and lifelong planning and learning.

Research in AI can be broadly categorized as follows:

- **Reasoning and problem solving:** the ability to imitate the step-by-step reasoning that humans use when they solve puzzles or make logical deductions;
- **Knowledge representation:** the ability to represent extensive knowledge about the world, including objects, properties, categories, and relations between these; situations, events, states, and time; causes and effects; and knowledge about knowledge;
- **Planning:** the ability to describe a set of possible actions that will yield a desired result, based on descriptions of the initial state, desired goals, and possible actions;
- **Learning:** the ability to improve automatically through experience;
- **Natural language processing:** the ability to read and understand human language;
- **Perception:** the ability to use input from sensors, such as cameras, microphones, tactile sensors, sonar, and others, to deduce aspects of the world; an example of perception is computer vision, which is the ability to analyze visual input; and
- **Motion and manipulation:** the ability to determine how to get from one point to another, and to execute that movement, including maintaining appropriate physical content with one or more objects; and
- **Social intelligence:** the ability to predict the actions of others by understanding their motivations and emotional states.

To advance the above fundamental research areas, NSF invests over \$100 million annually through a broad array of core as well as crosscutting programs. These investments span the “AI technology stack,” including real-time sensing and data acquisition, massive data management, machine learning, domain-specific modeling, AI infrastructure, human-AI interaction, and autonomy. Advances across the AI technology stack in turn contribute to innovations in a myriad of sectors.

For example, a team of NSF-funded computer scientists and pathologists at Stanford University has developed a model to teach computers to analyze breast cancer⁸. By assessing numerous novel morphological features of images of breast cancer tissue, the machine learning model can more accurately determine cancer diagnosis and prognosis than trained clinicians. A particularly striking revelation was that the cellular features that were the best predictors of patient survival were not from the cancer tissue itself, but rather from adjacent tissue, a discovery previously undetected by medical teams.

⁸ <http://stm.sciencemag.org/content/3/108/108ra113>

Similarly, an NSF-funded team of researchers at Johns Hopkins University have developed an AI program integrating data from the health records of more than 16,000 patients to identify 27 factors capable of predicting septic shock. Septic shock is a rapid immune response to infection that can cause organ failure, leading to more than 200,000 U.S. deaths annually. Early symptoms of septic shock are notoriously difficult to spot. But by combining and analyzing the 27 routine health factors – such as urine output and white-blood-cell counts – the program was able to accurately predict septic shock 85 percent of the time, usually before it harmed any organs.

NSF-funded researchers at the Texas Advanced Computing Center, the University of Texas Center for Transportation Research, and the City of Austin have also developed a new deep learning tool that uses raw camera footage from City of Austin cameras coupled with high-performance computing to recognize objects – people, cars, buses, trucks, bicycles, motorcycles, and traffic lights – and characterize how those objects move and interact⁹. This information can then be analyzed and queried by traffic engineers and officials to better determine traffic patterns and potential safety issues. For example, researchers were able to automatically identify a number of cases where vehicles and pedestrians were in close proximity, demonstrating how the system discovers dangerous locations without human intervention.

Machine learning approaches are increasingly used in security-sensitive applications such as malware detection and network intrusion detection. However, classical machine learning algorithms are not designed to operate in the presence of adversaries – and intelligent and adaptive adversaries may actively manipulate the information that they present in attempts to evade trained classifiers, leading to a cat-and-mouse game between the designers of learning systems and attackers who wish to evade them. A research team at the University of Virginia is actively developing automated techniques for predicting how well classifiers will resist adversaries' evasion attempts, along with general methods to automatically harden machine learning classifiers against adversarial evasion attacks¹⁰.

Big Data and Advanced Computing: Drivers of Modern AI

As noted earlier, advances in AI rely upon the availability of deep, high-quality, and accurate training datasets as well as advanced, scalable computing resources.

Indeed, we heard during the first AI hearing before this subcommittee several weeks ago that agencies across the Federal Government are being called upon to prioritize open training data and the release of federal data for this purpose^{11,12}. NSF is contributing to this effort, and to open science more broadly, through a variety of approaches. Since 2011, NSF has required all proposals to provide information about plans for data management and sharing of results of NSF-funded research. Prospective principal investigators must provide sufficient information to enable reviewers and program directors to assess both current plans and past performance. The research community, through merit review and program management, assesses the reasonableness of proposed data management and access. In 2015, NSF published a plan outlining a framework for activities to increase public access to scientific publications and digital scientific data resulting from NSF-funded research. The plan, *Today's Data, Tomorrow's*

⁹ <https://www.tacc.utexas.edu/-/artificial-intelligence-and-supercomputers-to-help-alleviate-urban-traffic-problems>

¹⁰ https://www.nsf.gov/awardsearch/showAward?AWD_ID=1619098

¹¹ <https://oversight.house.gov/hearing/game-changers-artificial-intelligence-part/>

¹² https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf

*Discoveries*¹³, sets forth the requirement that NSF-funded investigators are expected to share with other researchers, at no more than incremental cost and within a reasonable time, the primary data, samples, physical collections, and other supporting materials created or gathered in the course of work under NSF grants. Grantees are expected to encourage and facilitate such sharing. Pursuant to that plan, NSF also requires that articles derived from NSF-funded research appearing in peer-reviewed scholarly journals and juried conference proceedings or transactions be deposited in a public access-compliant repository and be available for download, reading, and analysis within one year of publication. In early 2017, NSF issued a “Request for Input on Federal Datasets with Potential to Advance Data Science¹⁴,” encouraging identification of possible datasets held by federal departments, agencies, and offices that would be useful in furthering research in machine learning and AI.

Beyond data, access to advanced, scalable tools and computational resources is critical to the success of NSF’s investments in AI R&D. NSF has long supported high-performance computing resources to accelerate fundamental science and engineering. Key NSF foci have included fundamental discoveries to support future generations of advanced computing; research and cyberinfrastructure promoting cohesive platforms and interoperability for large-scale data analytics as well as modeling and simulation; and support for a comprehensive advanced computing ecosystem for science and engineering research. Collectively, these foci emphasize a holistic approach to America’s science and engineering computational infrastructure, spanning both human and technical dimensions. More recently, NSF has partnered with commercial cloud providers, namely Amazon Web Services, Google Cloud Platform, and Microsoft Azure, to make available \$12 million in cloud resources to the academic research community over a three-year period^{15,16}. This public-private collaboration represents a powerful step forward in broadening access to the unique capabilities of the commercial cloud for data storage, processing, and analytics. It also allows NSF to balance various computational workloads across the full portfolio of advanced computing infrastructure.

AI Education and Workforce Development

NSF’s investments in AI research are accompanied by investments in education and workforce development. Research undertaken in academia not only engages some of our Nation’s best and brightest researchers, but because these researchers are also teachers, new generations of students are exposed to the latest thinking from the people who understand it best. Further, as these students graduate and transition into the workplace, they bring this knowledge and understanding with them.

NSF is funding research and development that is building the necessary foundations for implementing rigorous and engaging computer science (CS) education at all levels: preK-12, colleges/universities, and continuing education programs. These investments are predicated on the importance of a diverse workforce that understands foundational concepts of CS and computational thinking, knows how to creatively use and develop new methodologies and tools including in AI, has the capacity to interact with all sectors of our society, and is prepared to lead the global information economy.

¹³ <https://www.nsf.gov/pubs/2015/nsf15052/nsf15052.pdf>

¹⁴ <https://www.nsf.gov/pubs/2017/nsf17049/nsf17049.jsp>

¹⁵ https://www.nsf.gov/news/news_summ.jsp?cntn_id=244450

¹⁶ <https://www.microsoft.com/en-us/research/blog/microsoft-boosts-nsfs-big-data-regional-innovation-hubs-3m-cloud-computing-credits/>

Over the last decade, NSF investments have laid the groundwork for rigorous and engaging computer science education at the preK-12 level for all students in all parts of the U.S.. For example, sustained NSF support for nearly a decade has enabled the development of a new Advanced Placement® (AP®) Computer Sciences Principles (CSP) framework and exam, along with a number of aligned curricula and associated teacher professional development resources. The first official AP CSP exam was administered in May 2017, and proved to be the largest launch of any AP course in the 60-year history of The College Board: over 2,500 schools offered AP CSP courses, and they combined to enable more than 50,000 students to take the exam. Additionally, compared to participation in the existing AP computer science exam (CS A), African American participation was 7% in CSP (versus 4% in CS A); Hispanic participation was 19% (versus 11%), and female participation was 30% (versus 25%). Today, NSF is continuing to grow the knowledge base and capacity for high-quality preK-12 computer science education, as well as scalable and sustainable models of professional development for educators. In achieving these results at the preK-12 level, NSF has partnered with other federal agencies including the U.S. Department of Education and the U.S. Department of Defense. In addition, NSF has worked with private partners, resulting in a number of NSF projects being scaled nationally.

NSF has also made significant investments at the undergraduate level, where computer science departments are experiencing a surge of non-major students in their mid-level and advanced courses. Over the last several years, NSF has sought to support computer science departments and universities in responding to this changing landscape, restructuring departments and universities to better prepare students to employ the power of computing across the interdisciplinary and multidisciplinary collaborations of the future. A key focus has been on “CS+X” – with “X” constituting another discipline, sector, or societal grand challenge – and the dynamic needs of many industry sectors.

These efforts at the undergraduate level have been complemented by NSF’s longstanding support for Research Experiences for Undergraduates (REU) Sites¹⁷ in the area of AI. REU Sites are based on independent proposals that seek to initiate and conduct projects engaging a number of undergraduate students in research. Each REU Site must have a well-defined common focus, based in a single discipline or spanning interdisciplinary or multidisciplinary research opportunities with a coherent intellectual theme, which enables a cohort experience for participating students. NSF recently awarded an REU Site to Carnegie Mellon University that is providing high-quality guided research experiences for undergraduate students in computer vision, field and space robotics, artificial intelligence, manipulation, and machine learning¹⁸.

NSF supports graduate students through research grants to individual investigators, as well as through the Graduate Research Fellowships (GRF) program. Beginning in January 2018, in the evaluation of GRF applications submitted to NSF, applicants with “computationally- or data-intensive” research plans are identified as a priority. This focus on computational or data-intensive research spans applicants across all of NSF. In fiscal year (FY) 2017, more than 6,500 graduate students in computer and information science and engineering (and closely related fields) were supported on NSF awards.

¹⁷ https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=5517

¹⁸ https://www.nsf.gov/awardsearch/showAward?AWD_ID=1659774

NSF has also funded early-career investigators through the Faculty Early-Career Development (CAREER) program¹⁹, which offers NSF's most prestigious research award in support of early-career faculty. One active award at the University of California, San Diego is supporting research to develop private and highly efficient machine learning tools for classification and clustering of patient medical records²⁰. This research could help lead to the discovery of population-wide patterns enabling advances in genetics, disease mechanisms, drug discovery, healthcare policy, and public health.

NSF's Cyberlearning for Work at the Human-Technology Frontier program²¹ is seeking to respond to the pressing need to educate and re-educate learners of all ages (students, teachers, and workers) to ultimately function in highly technological environments, including in collaboration with intelligent systems. An important direction of this program is to foster lifelong learning with and through technology, particularly in preparation for and within the context of the work setting. This program invites transformative proposals that integrate advances in what is known about how people learn (individually and in groups) with the opportunities offered by new and emerging technologies such as AI to prepare future learners and workers.

NSF's Ten Big Ideas: Harnessing Data for 21st Century Science and Engineering and The Future of Work at the Human-Technology Frontier

In FY 2016, NSF announced a set of bold questions that will drive the agency's long-term research agenda – questions that will ensure future generations continue to reap the benefits of fundamental research. These 10 “Big Ideas” aim to capitalize on what NSF does best: catalyze interest and investment in fundamental research, which is the basis for discovery, invention, and innovation, along with education. The Big Ideas define a set of cutting-edge research agendas and processes that are suited for NSF's broad portfolio of investments, and will require collaborations with industry, private foundations, other agencies, science academies and societies, and universities. These ideas will push forward the frontiers of US research and provide innovative approaches to solve some of the most pressing problems the world faces, as well as lead to discoveries not yet known. They will provide platforms to bring together every field of study, from science and education, to engineering and astrophysics, to radically alter the conduct of science and engineering across the scientific enterprise in a manner that is not possible by simply continuing discipline-specific efforts at current levels.

Two of the 10 Big Ideas have a strong AI focus. First, Harnessing the Data Revolution for 21st-Century Science and Engineering engages NSF's research community in the pursuit of fundamental research in data science and engineering; the development of a cohesive, federated, national-scale approach to research data infrastructure; and the development of a 21st-century data-capable workforce. Second, The Future of Work at the Human-Technology Frontier aims to help build an understanding of how constantly evolving technologies are actively shaping the lives of workers and how people in turn can shape those technologies, especially in the world of work. This Big Idea will bring together NSF research communities to conduct fundamental scientific research on the interaction of humans, society, and technology that will help shape the future of work to increase opportunities for workers and productivity for the American economy.

¹⁹ https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503214

²⁰ https://www.nsf.gov/awardsearch/showAward?AWD_ID=1253942

²¹ https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504984

Engaging with Industry: Transitioning Research Innovations to Practice

In addition to the partnership with leading cloud providers mentioned previously, NSF has several focused programs to create and expand partnerships with the business community. The Industry-University Cooperative Research Centers (IUCRC) program was created in 1973 to develop long-term partnerships among industry, academia, and government. NSF invests in these partnerships to promote research of mutual interest, contribute to the Nation's research infrastructure base, enhance the intellectual capacity of the engineering and science workforce, and facilitate technology transfer. NSF currently supports 77 IUCRCs involving over 200 university sites. Each center has, on average, approximately 17 industrial partners. For every dollar provided to a center from the NSF IUCRC program, approximately seven dollars are provided by the industry members and other sources. More than 2,000 students conduct research at IUCRCs each year, and approximately 30% of those students graduating each year are hired by the center's member companies. For example, the NSF IUCRC for Big Learning²², located at the University of Florida and the University of Missouri at Kansas, seeks to create state-of-the-art deep learning methodologies and technologies and enable intelligent applications. This IUCRC is pursuing deep learning algorithms in embedded systems for mobile and Internet-of-Things (IoT) applications in a number of domains, including healthcare and cybersecurity.

NSF also provides a much-needed bridge across the so-called "valley of death" between R&D and commercialization. NSF-funded AI research has led to the formation of numerous start-up companies, enabling transition of research results to deployment and implementation. NSF has supported these start-ups through its Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs, which stimulate technological innovation in the private sector by strengthening the role of small business concerns in meeting federal R&D needs. Similarly, the NSF Innovation Corps™ (NSF I-Corps™) program has provided AI grantees with entrepreneurial education so that they can conduct customer discovery and successfully pursue valuable product opportunities that emerge from their academic research. Since the inception of the NSF I-Corps program in 2011, several I-Corps Teams in the AI domain have participated in the curriculum. For example, an I-Corps project at the Massachusetts Institute of Technology is investigating the commercial applications of a machine learning system that ingests, aggregates, processes, and learns mappings about buildings and their environments from multiple data sources²³. This system would allow users to make inferences about built and energy infrastructure and their surroundings where data are not available in a low-cost, large-scale, automated way, removing the need for certain expensive sensor deployments or manual surveys.

Finally, NSF is constantly exploring new models of partnership with industry. Over the last decade, NSF has increasingly partnered with industry, including with the Semiconductor Research Corporation, Intel Labs University Collaboration Office, and VMware, Inc., on programs that jointly fund research projects advancing the state of the art in specific areas. Additionally, in 2016, NSF helped catalyze the formation of a new industry consortium comprising over 25 companies and associations in the wireless networking sector to support the design, development, deployment, and initial operations of city-scale testing platforms that will accelerate fundamental research on wireless communication and networking technologies. This collaboration combines \$50 million from NSF over seven years with \$50 million in cash and in-kind contributions from the industry consortium to advance wireless technologies beyond fifth-generation ("5G") networks. This new public-private partnership serves as a compelling model for potential future government-industry collaborations in AI R&D.

²² https://www.nsf.gov/awardsearch/showAward?AWD_ID=1747751;

https://www.nsf.gov/awardsearch/showAward?AWD_ID=1747783

²³ https://www.nsf.gov/awardsearch/showAward?AWD_ID=1820773

Coordination and Collaboration Across the Federal Government

NSF's close coordination and collaboration with other federal agencies pursuing AI R&D is another critically important factor in shaping its long-term investments. NSF serves as co-chair of the Networking and Information Technology Research and Development (NITRD) Subcommittee of the NSTC, which coordinates all investments in fundamental networking and information technology R&D across more than 20 member departments, agencies, and offices of the Federal Government²⁴. The full NITRD portfolio spans more than \$4 billion annually. As part of the NITRD program, NSF co-chairs a number of interagency working groups, including one focused on Intelligent Robotics and Autonomous Systems (IRAS)²⁵.

Since spring 2016, NSF has actively participated in the NSTC MLAI Subcommittee, and co-chaired the *National Artificial Intelligence Research and Development Strategic Plan* that the subcommittee published in 2016²⁶. NSF's investments in AI are also strongly aligned with the FY 2019 Administration R&D Budget Priorities^{27,28}, and the *National Security Strategy of the United States of America* published in December 2017²⁹.

Renewing NSF

Apart from its support of extramural research at the frontiers of science and engineering, NSF also strives to be a model federal agency operationally. Given that the landscape in which NSF executes its mission is constantly evolving – for example, research questions are becoming increasingly interdisciplinary in nature, and thus they require new levels and forms of scientific and engineering collaboration – NSF constantly pursues organizational reforms to optimize its efficiency and effectiveness, including its customer service standards.

NSF's current set of operational reforms, collectively called "Renewing NSF," includes a focus on *Making Information Technology (IT) Work for Us*. This focus area seeks to accelerate modernization of NSF's IT infrastructure via adoption of cloud offerings, consolidated computing platforms, software-defined network infrastructure, and automated change management processes to improve overall resilience of NSF's systems. It also aims to promulgate adoption of automated, intelligent tools to substantively evolve NSF's business processes, including merit review of submitted proposals. For example, beginning in FY 2018, NSF is conducting a pilot in two of its divisions to leverage AI technologies, including natural language processing and machine learning, to automate the selection of peer reviewers for proposals – traditionally a time-intensive and manual process – based on the contents of submitted proposals.

Ultimately, NSF strives to implement leading-edge IT solutions that can be adapted easily and quickly, and that enhance employee productivity and satisfaction by enabling access to readily available, reliable, and fully integrated data to support decision making.

²⁴ <http://www.nitrd.gov/>

²⁵

[https://www.nitrd.gov/nitrdgroups/index.php?title=Subcommittee_on_Networking_and_Information_Technology_Research_and_Development_\(NITRD_Subcommittee\)](https://www.nitrd.gov/nitrdgroups/index.php?title=Subcommittee_on_Networking_and_Information_Technology_Research_and_Development_(NITRD_Subcommittee))

²⁶ https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf

²⁷ <https://www.whitehouse.gov/sites/whitehouse.gov/files/ostp/fy2019-administration-research-development-budget-priorities.pdf>

²⁸ https://www.whitehouse.gov/wp-content/uploads/2018/02/ap_18_research-fy2019.pdf

²⁹ <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>

Conclusions

NSF has made significant investments in foundational and multidisciplinary AI research over the last several decades. These investments have enhanced our Nation's economic competitiveness and national security, in direct alignment with NSF's mission "to promote the progress of science; to advance the national health, prosperity, and welfare; and to secure the national defense." They have also given rise to fundamentally new research directions and opportunities for the future. NSF's interdisciplinary education research portfolios are contributing to a next-generation workforce capable of pursuing AI research and taking on new jobs that will soon be created across multiple sectors of the economy. Across our research and education investments, NSF partnerships with other federal agencies and industry are also helping to advance machine learning and AI, and transition innovations into the marketplace.

The development of AI is advancing at a rapid pace, and NSF will continue to invest in fundamental AI research, infrastructure, and workforce development to maintain U.S. global leadership in this field. With sustained support for AI R&D in both the executive and legislative branches, there is a unique opportunity to generate breakthroughs that can further our national priorities. AI holds the potential to transform the lives of Americans through increased economic prosperity, improved educational opportunities and quality of life, and enhanced national and homeland security. A just-published strategic plan by the U.S. Government Accountability Office describing *Trends Affecting Government and Society*³⁰ identified AI and automation as one of five emerging technologies that will potentially transform society and noted, "The extent to which the United States is able to focus R&D investment in key technology areas will be a key factor in U.S. competitiveness in the global economy."

Allow me to leave you with four key characteristics about AI:

- **AI is broad:** it involves many areas of study, including decision making, natural language processing, machine learning, and robotics, and is closely allied with other fields of computer and information science and engineering, including data science, algorithms, and advanced computing systems;
- **AI is impactful:** it has strong impacts across all application domains, ranging from precision medicine to energy to transportation to education and learning;
- **AI is challenging:** it offers many open R&D challenges, and these provide rich intellectual domains of study with many long-term challenges; and
- **AI is growing:** it is engaging ever-growing numbers of students and faculty who are joining the AI R&D community.

Thank you for the opportunity to testify before the Information Technology Subcommittee on this very important and timely topic. I would be happy to answer any questions at this time.

³⁰ <https://www.gao.gov/assets/700/690262.pdf>