2016–2019 PROGRESS REPORT: ADVANCING ARTIFICIAL INTELLIGENCE R&D

A report by the

ARTIFICIAL INTELLIGENCE RESEARCH & DEVELOPMENT INTERAGENCY WORKING GROUP

SUBCOMMITTEE ON NETWORKING & INFORMATION TECHNOLOGY RESEARCH & DEVELOPMENT

SUBCOMMITTEE ON MACHINE LEARNING & ARTIFICIAL INTELLIGENCE

and the SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE

of the

NATIONAL SCIENCE & TECHNOLOGY COUNCIL

NOVEMBER 2019

About the National Science and Technology Council

The National Science and Technology Council (NSTC) is the principal means by which the Executive Branch coordinates science and technology policy across the diverse entities that make up the Federal research and development enterprise. A primary objective of the NSTC is to ensure that science and technology policy decisions and programs are consistent with the President's stated goals. The NSTC prepares research and development strategies that are coordinated across Federal agencies aimed at accomplishing multiple national goals. The work of the NSTC is organized under committees that oversee subcommittees and working groups focused on different aspects of science and technology. More information is available at https://www.whitehouse.gov/ostp/nstc.

About the Office of Science and Technology Policy

The Office of Science and Technology Policy (OSTP) was established by the National Science and Technology Policy, Organization, and Priorities Act of 1976 to provide the President and others within the Executive Office of the President with advice on the scientific, engineering, and technological aspects of the economy, national security, homeland security, health, foreign relations, the environment, and the technological recovery and use of resources, among other topics. OSTP leads interagency science and technology policy coordination efforts, assists the Office of Management and Budget with an annual review and analysis of Federal research and development (R&D) in budgets, and serves as a source of scientific and technological analysis and judgment for the President with respect to major policies, plans, and programs of the Federal Government. More information is available at https://www.whitehouse.gov/ostp.

About the Select Committee on Artificial Intelligence

The Select Committee on Artificial Intelligence (AI) advises and assists the NSTC to improve the overall effectiveness and productivity of Federal R&D efforts related to AI to ensure continued U.S. leadership in this field. It addresses national and international policy matters that cut across agency boundaries, and it provides formal mechanisms for interagency policy coordination and development for Federal AI R&D activities, including those related to autonomous systems, biometric identification, computer vision, human-computer interactions, machine learning, natural language processing, and robotics. It also advises the Executive Office of the President on interagency AI R&D priorities; works to create balanced and comprehensive AI R&D programs and partnerships; leverages Federal data and computational resources across department and agency missions; and supports a technical, national AI workforce.

About the Subcommittee on Machine Learning and Artificial Intelligence

The Machine Learning and Artificial Intelligence (MLAI) Subcommittee monitors the state of the art in machine learning (ML) and artificial intelligence within the Federal Government, in the private sector, and internationally to watch for the arrival of important technology milestones in the development of AI, to coordinate the use of and foster the sharing of knowledge and best practices about ML and AI by the Federal Government, and to consult in the development of Federal MLAI R&D priorities. The MLAI Subcommittee reports to the Committee on Technology and the Select Committee on AI. The MLAI Subcommittee also coordinates AI taskings with the Artificial Intelligence Research and Development Interagency Working Group (see below).

About the Subcommittee on Networking & Information Technology Research & Development

The Networking and Information Technology Research and Development (NITRD) Program is the Nation's primary source of Federally funded work on pioneering information technologies (IT) in computing, networking, and software. The NITRD Subcommittee guides the multiagency NITRD Program in its work to provide the R&D foundations for ensuring continued U.S. technological leadership and meeting the needs of the Nation for advanced IT. It reports to the NSTC Committee on Science and Technology Enterprise. The Subcommittee is supported by the interagency working groups (IWGs) that report to it and by its National Coordination Office. More information is available at https://www.nitrd.gov/about/.

About the Artificial Intelligence Research & Development Interagency Working Group

The NITRD AI R&D Interagency Working Group coordinates Federal R&D in AI; it also supports and coordinates activities tasked by the Select Committee on AI and the NSTC Subcommittee on Machine Learning and Artificial Intelligence. This vital work promotes U.S. leadership and global competitiveness in AI R&D. The AI R&D IWG spearheaded the 2019 update of the National Artificial Intelligence Research and Development Strategic Plan. More information is available at https://www.nitrd.gov/groups/AI.

About this Document

This document reports current trends and examples of Federal R&D investments, program information, and activities in artificial intelligence that directly address the R&D challenges and opportunities noted in *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, available at https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf.

Copyright Information

This document is a work of the United States Government and is in the public domain (see 17 U.S.C. §105). It may be freely distributed, copied, and translated, with acknowledgment to OSTP; requests to use any images must be made to OSTP. Published in the United States of America, 2019.

National Science & Technology Council

Chair

Kelvin Droegemeier, Director, OSTP

Staff

Chloé Kontos, Executive Director, NSTC

Select Committee on Artificial Intelligence

Co-Chairs

France A. Córdova, Director, NSF

Steven Walker, Director, DARPA

Michael Kratsios, Chief Technology Officer of the United States

Subcommittee on Machine Learning and Artificial Intelligence

Co-Chairs

Lynne Parker, Assistant Director for Artificial Intelligence, OSTP

Erwin Gianchandani, Assistant Director (Acting), CISE, NSF

Charles Romine, Director, Information Technology Laboratory, NIST

Stephen Binkley, Deputy Director, Science Programs, Office of Science, DOE

Executive Secretary

Faisal D'Souza, NITRD NCO

Subcommittee on Networking & Information Technology Research & Development

Co-Chairs

Kamie Roberts, Director, NITRD NCO

Erwin Gianchandani, Assistant Director (Acting), CISE, NSF

Executive Secretary

Nekeia Butler, NITRD NCO

Artificial Intelligence Research & Development Interagency Working Group

Co-Chairs

Jeff Alstott, Program Manager, IARPA, Office of the Director of National Intelligence

Henry Kautz, Division Director, CISE Division of Information and Intelligent Systems, NSF

Staff

Faisal D'Souza, Technical Coordinator, NITRD NCO

Writing Team

Jeff Alstott, IARPA
Gil Alterovitz, VA
Sameer Antani, NIH
Charlotte Baer, NIFA
Kimberly Ferguson-Walter, U.S. Navy
William Ford, NIJ
Timothy Gallaudet, NOAA
Michael Garris, NIST
Erwin Gianchandani, NSF
Travis Hall, NTIA
Henry Kautz, NSF
Erin Kenneally, DHS
Dai H. Kim, OSD
David Kuehn, DOT
James Lawton, AFOSR
Steven Lee, DOE
Aaron Mannes, DHS
Anup Mathur, Census Bureau
Nikunj Oza, NASA
Dinesh Patwardhan, FDA
William Scherlis, DARPA
Pauline Smith, U.S. Army
Elham Tabassi, NIST

Table of Contents

Executive Summary

The United States national strategy for artificial intelligence (AI), the American AI Initiative, identifies research and development (R&D) as a top priority for maintaining global leadership in AI.[1] The United States leads the world in AI innovation, due in large part to its robust R&D ecosystem. Federal agencies contribute significantly to AI innovation by investing in numerous world-class research programs in areas consistent with the unique missions of each agency.

This *2016–2019 Progress Report on Advancing Artificial Intelligence R&D* ("2016–2019 Progress Report") documents the important progress that agencies are making to deliver on Federal AI R&D.

Guiding Federal research investments is the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* ("2019 Plan"),[2] which builds upon the 2016 version of the Plan. The 2019 Plan articulates eight national AI R&D strategies:

Strategy 1: Make long-term investments in AI research.

Strategy 2: Develop effective methods for human-AI collaboration.

Strategy 3: Understand and address the ethical, legal, and societal implications of AI.

Strategy 4: Ensure the safety and security of AI systems.

Strategy 5: Develop shared public datasets and environments for AI training and testing.

Strategy 6: Measure and evaluate AI technologies through benchmarks and standards.

Strategy 7: Better understand the national AI R&D workforce needs.

Strategy 8: Expand public-private partnerships in AI to accelerate advances in AI.

This 2016–2019 Progress Report highlights AI research first by strategy, then by sector, with subsequent supporting details describing individual agency contributions that provide a whole-of-government overview. The diversity of programs and activities reflects the remarkable breadth and depth of Federal investments in AI. This report highlights not only the broad themes of Federal R&D but also provides illustrative examples in sidebars that highlight individual agency AI R&D breakthroughs that advance the state of the field.

Taken as a whole, the 2016–2019 Progress Report conveys the following key messages:

1. The Federal Government is investing in a considerable breadth and depth of innovative AI concepts that can transform the state of the field.
2. The United States benefits significantly from the broad spectrum of Federal agencies that invest in AI from their unique mission perspectives, consistent with the national AI R&D strategy.
3. Federal investments have generated impactful breakthroughs that are revolutionizing our society for the better.

Collectively, the investments described in this report demonstrate how the Federal Government leverages and improves America's AI capabilities through R&D and ensures that those capabilities are increasing prosperity, safety, security, and quality of life for the American people for decades to come.

[1] https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/

[2] https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf

Abbreviations

AFOSR	Air Force Office of Scientific Research
AI	artificial intelligence
AIST	Advanced Information Systems Technology (NASA)
CISE	Computer and Information Science and Engineering Directorate (NSF)
co-robot	collaborative robot
DARPA	Defense Advanced Research Projects Agency
DATC	Data Analytics Technology Center (DHS)
DHS	Department of Homeland Security
DIVA	Deep Intermodal Video Analytics (IARPA)
DOC	Department of Commerce
DoD	Department of Defense
DOE	Department of Energy
DOT	Department of Transportation
EAR	Exploratory Advanced Research (DOT)
FACT	Food and Agriculture Cyberinformatics and Tools (NIFA)
FAIR	findable, accessible, interoperable, reusable
FARSAIT	Fundamental and Applied Research and Standards for AI Technologies (NIST)
FBI	Federal Bureau of Investigation
FDA	Food and Drug Administration
FHWA	Federal Highway Administration (DOT)
FW-HTF	Future of Work at the Human-Technology Frontier (NSF)
GSA	General Services Administration
HDR	Harnessing the Data Revolution (NSF)
HHS	Department of Health and Human Services
IARPA	Intelligence Advanced Research Projects Activity
IT	information technology/ies
JAIC	Joint Artificial Intelligence Center (DoD)
JARVIS DFT	Joint Automated [data] Repository for Various Integrated Simulations, Density Functional Theory (NIST)
ML	machine learning
MLAI	Machine Learning and Artificial Intelligence Subcommittee of the NSTC Committee on Technology
NASA	National Aeronautics and Space Administration
NCCoE	National Cybersecurity Center of Excellence (NIST)
NCO	National Coordination Office
NEX-AI	NASA Earth eXchange AI database
NIFA	National Institute of Food and Agriculture (USDA)
NIH	National Institutes of Health (HHS)
NIJ	National Institute of Justice
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development (Program)
NOAA	National Oceanic and Atmospheric Administration
NRI	National Robotics Initiative
NSF	National Science Foundation
NSTC	National Science and Technology Council
NTIA	National Telecommunications and Information Administration
ONR	Office of Naval Research
OSD	Office of the Secretary of Defense
OSTP	Office of Science and Technology Policy
R&D	research and development
SBE	Social, Behavioral, and Economic Sciences Directorate (NSF)
SICLF	Science of Information, Computation, Learning, and Fusion (AFOSR)
STEM	science, technology, engineering, and mathematics
STRIDES	Science and Technology Research Infrastructure for Discovery, Experimentation & Sustainability (NIH)
TrojAI	Trojans in AI (IARPA)
USDA	U.S. Department of Agriculture
VA	Department of Veteran Affairs
XAI	Explainable Artificial Intelligence (DARPA)

Introduction

Global leadership in AI R&D is a national imperative highlighted by the President's February 2019 Executive Order on Maintaining American Leadership in Artificial Intelligence.[3] This *2016–2019 Progress Report on Advancing Artificial Intelligence R&D* ("2016–2019 Progress Report") provides an update on Federal AI R&D activities over the last three years, with specific examples of recent and current Federal agency programs and activities aligned with the eight strategies described in *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*.[4] This 2016–2019 Progress Report illustrates the remarkable breadth, diversity, impact, and importance of Federal investments in AI R&D. The activities listed here are highlights that are responsive to and well-aligned with the national AI R&D strategic plan. They offer a glimpse into how Federal AI R&D activities are driving advances in AI and in turn impacting many sectors of the economy. While this 2016–2019 Progress Report is not meant to be exhaustive, it is reflective of the significant cross-agency activities supported by the Federal Government. Additional discussions of agency R&D activities can be found via the AI.gov website.[5] Additionally, a variety of smaller, cross-sector, or classified activities unable to be noted here significantly contribute to meeting the Nation's commitment to responsibly advancing AI.

National AI R&D Strategy

In June 2019, the Administration released *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* ("2019 Plan"). The 2019 Plan identifies priority areas for Federal investments in AI R&D to ensure that the Federal Government leverages and improves America's AI capabilities, and that those capabilities increase prosperity, safety, security, and quality of life for the American people for decades to come. *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* built on the 2016 *National AI R&D Strategic Plan*, taking into account the new research, technical innovations, and other considerations that have emerged over the three years since the release of the 2016 Plan. The eight national AI R&D strategies articulated in the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* are:

Strategy 1: Make long-term investments in AI research.

Strategy 2: Develop effective methods for human-AI collaboration.

Strategy 3: Understand and address the ethical, legal, and societal implications of AI.

Strategy 4: Ensure the safety and security of AI systems.

Strategy 5: Develop shared public datasets and environments for AI training and testing.

Strategy 6: Measure and evaluate AI technologies through benchmarks and standards.

Strategy 7: Better understand the national AI R&D workforce needs.

Strategy 8: Expand public-private partnerships in AI to accelerate advances in AI.

This 2016–2019 Progress Report is structured around these eight strategies. For each strategy, agency programs and activities that address that strategy are highlighted, with an emphasis on efforts undertaken over the last three years, subsequent to the release of the 2016 *National AI R&D Strategic Plan*. Table 1 summarizes the involvement of Federal R&D agencies in each of the eight strategies; these involvements are explained in more detail in the remainder of the report. The overall picture that

[3] Executive Order 13859: https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/

[4] https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf

[5] https://www.whitehouse.gov/ai/ai-american-innovation/

emerges is that of a comprehensive set of investments, progress, and accomplishments, along with noteworthy continuity in Federal AI R&D strategy across multiple Administrations.

The 2016–2019 Progress Report will be updated periodically to document the Nation's progress in meeting the aims of the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*. In addition, starting with the FY2020 document[6] the annual *NITRD Supplement to the President's Budget* publishes overall projected Federal AI R&D budget figures by agency (not including the AI R&D programs of the Department of Defense and the Defense Advanced Research Projects Agency).

Table 1. Summary of Investments by Federal Entities in the Eight National AI R&D Strategies

AI R&D Strategies	AFOSR	Army	Census	DARPA	DHS	DoD*	DOE	DOT	FBI	FDA	GSA	HHS	IARPA	NASA	NIFA	NIH	NIJ	NIST	NOAA	NSF	NTIA	ONR	VA
1. Make long-term investments in AI research	X	X		X	X	X	X	X	X	X	X		X	X	X	X		X	X	X	X	X	
2. Develop effective methods for human-AI collaboration	X	X		X	X	X	X	X		X			X		X	X	X	X	X	X		X	
3. Understand and address the ethical, legal, and societal implications of AI		X		X	X	X	X	X							X	X				X	X		
4. Ensure safety and security of AI systems		X		X		X	X			X			X	X				X	X	X			
5. Develop shared public datasets and environments for AI training and testing						X	X	X		X	X	X	X	X		X		X	X	X	X	X	X
6. Measure and evaluate AI technologies through benchmarks and standards				X	X	X	X			X			X			X	X	X	X			X	
7. Better understand the national AI R&D workforce needs			X			X	X						X		X		X	X	X	X			
8. Expand public-private partnerships in AI to accelerate advances in AI		X		X	X	X	X		X	X		X			X	X	X	X	X	X			

***** DoD performs R&D in the AI R&D Strategies through coordinated activities by its service agencies.

How to Read this Report

To help the reader explore topical themes, agency investments are broadly organized around the topical areas of science and engineering, defense, the economy, health, and justice and security. To quickly understand the importance and impact of the R&D investments, the casual reader is encouraged to focus on the sidebars, which describe particularly significant research and program results from Federal agency AI R&D programs.

[6] https://www.nitrd.gov/pubs/FY2020-NITRD-Supplement.pdf

Strategy 1: Make long-term investments in AI research

The past several years have seen dramatic advances in the ability of AI systems to solve real-world problems of critical national importance. Scientists supported by, or directly working for, Federal agencies have developed AI applications ranging from medical diagnosis to homeland security. The most prominent technical advances have been in methods and applications of machine learning (ML) using statistical models or artificial neural networks. ML is most readily applied to problems that require recognizing patterns, classifying entities, or predicting values where large amounts of ground-truth labeled data are available for training the system. As noted in the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, long-term research investments must be made in ML on such topics as reducing the amount of labeled data needed to train a model, interactive and continuous learning, mitigating unintended biases in models, and increasing algorithmic robustness in the presence of adversaries attempting to trick the system into misclassifying inputs—a vital issue when ML is applied to problems of security.

As important as it is, ML is not the whole story of AI. Agencies recognize the importance of supporting long-term research in other areas of artificial intelligence as well as in design for comprehensive reasoning architectures. These other AI areas include representing and reasoning with commonsense background knowledge, decision making and decision support, principles and algorithms for problem solving by teams of humans and AI systems, robotics or embodied AI, and methods for quickly finding the best solution to a problem among an exponential number of options. This last task, called combinatorial optimization, is important for problems in planning, scheduling, and resource allocation.

The National Science Foundation (NSF) has supported basic research in AI since the earliest days of the field; indeed, many of today's commercial AI applications are founded on the results of decades of NSF-funded research. NSF continues to fund foundational research in AI through a range of programs. For example, the Division of Information and Intelligent Systems (IIS) within the Directorate for Computer and Information Science and Engineering (CISE) supports all areas of basic AI research in its longstanding Robust Intelligence and Information Integration and Informatics programs. A growing number of cross-directorate programs also fund foundational as well as use-inspired AI research. One of the largest such programs is the National Robotics Initiative (NRI) 2.0, a joint program between the NSF Directorates for CISE; Engineering; Education and Human Resources; and Social, Behavioral and Economic Sciences (SBE); and a number of agency partners, including the Air Force Office of Scientific Research (AFOSR), Department of Defense (DoD), Department of Energy (DOE), National Aeronautics and Space Administration (NASA), Office of Naval Research (ONR), and U.S. Department of Agriculture National Institute of Food and Agriculture (USDA/NIFA). NRI 2.0 focuses on fundamental research on collaborative robots (co-robots) that work beside or cooperatively with people. Another example is NSF's National Artificial Intelligence Research Institutes program, announced in October 2019, that will support advancement of multidisciplinary, multistakeholder research on larger-scale, longer-time-horizon challenges in AI research than are supported in typical Federal research grants. Through this program, which involves collaboration with several other Federal agency partners, each institute will receive funding for up to five years to advance foundational as well as translational AI research.

Science & Engineering

AI benefits from advances in data science, which lies at the intersection of computer science, mathematics, and statistics and includes methods for collecting and managing data, as well as applications of these technologies to problems in science and engineering. Data science has been

the focus of the cross-directorate NSF program Critical Techniques, Technologies, and Methodologies for Advancing Foundations and Applications of Big Data Sciences and Engineering (BIGDATA) and a series of new programs announced in 2019 under Harnessing the Data Revolution (HDR), one of NSF's ten Big Ideas for Future Investment. HDR will support data science institutes across the Nation through a two-phase competitive process. HDR will also expand an existing program, Transdisciplinary Research in Principles of Data Science (TRIPODS), which supports research in the foundations of statistics and machine learning by bringing together computer scientists, mathematicians, statisticians, and electrical engineers.

Foundational research that links AI to neuroscience is supported by the NSF-led international program on Collaborative Research in Computational Neuroscience. Research that advances both AI and medicine is supported by the joint NSF–National Institutes of Health (NIH) program on Smart and Connected Health. The Advances in Biological Informatics program run by the NSF Directorate for Biological Sciences supports many projects that advance the state of the art of machine learning and knowledge representation for biological domains.

Across DOE, ML is transforming science and energy research. In January 2018, the DOE Basic Research Needs workshop established the following priority research directions for AI R&D:[8] (1) domain-aware scientific machine learning for leveraging scientific domain knowledge; (2) interpretable scientific machine learning for explainable and understandable results; (3) robust scientific machine learning for stable, well-posed, and efficient methods; (4) data-intensive scientific machine learning for automated

NSF	*A New Way to Count*

Counting seems simple, but hard counting problems lie at the heart of many problems in artificial intelligence as well as other areas of science and engineering. In particular, probabilistic reasoning—determining the likelihood of a particular outcome given a knowledge base—is equivalent to counting the number of possible states of the world that are consistent with both the outcome and the knowledge base. All known algorithms for performing an exact count require runtime that grows exponentially with the size of the knowledge base, far too long for real-world problems, even with the fastest computers. AI researchers supported by NSF's Expeditions in Computing program[7] developed a clever and fast method to compute a good estimate of the count.

To get a sense of how the algorithm works, consider counting the number of people in a large auditorium. Suppose you tell everyone to stand; then ask each person to flip a coin, and if it comes up heads, they should sit down. Repeat the cycles of flipping and sitting until everyone is seated. One can show, with high probability, that the number of people is the number "two" raised to the power of the number of cycles. The algorithm the researchers invented adds random constraints to the knowledge base—simulating the flipping of coins—until the knowledge base is inconsistent. Further techniques using hashing and optimization are employed by the algorithm to make it robust in practice.

This new counting algorithm not only makes it possible to solve certain problems in probabilistic reasoning that could not be solved before, but it also has found use in materials science, computer vision, and software verification. Perhaps surprisingly, this discovery came from an NSF award that was focused on using AI to design new materials for batteries for improved environmental and economic sustainability. This project provides an excellent example of the interplay between fundamental and applied research, and how the creation of new general AI algorithms can advance science and engineering in unexpected areas far beyond the original source of motivation.

[7] NSF Award #1522054, Expanding the Horizons of Computational Sustainability.

[8] https://coefs.uncc.edu/tschmit4/files/2019/02/doe_machine_learning.pdf

scientific inference and data analysis; (5) machine learning-enhanced modeling and simulation for predictive scientific computing; and (6) intelligent automation and decision support for the management and control of complex systems, including for improving the operational capabilities of scientific user facilities, communication networks, power grids, and other sensor-equipped infrastructures and complex processes.

Artificial Intelligence is also a strategic priority for the National Institute of Standards and Technology (NIST). Through targeted and sustained investments, NIST has been working to expand its efforts to address emerging research and standards challenges surrounding AI to build the confidence and trust that is needed to accelerate technology adoption and deployment and create an expanded commercial marketplace. As a trusted, rigorous, nonregulatory research organization, NIST is well positioned to deliver tools to measure and understand the capabilities and limitations of AI technologies and the underlying data that catalyze them. NIST's principal role in AI is to develop public trust and confidence in AI systems. In FY2018, NIST initiated the Fundamental and Applied Research and Standards for AI Technologies (FARSAIT) program. This program is focused on enhancing trust in AI systems by measuring and improving the security, explainability, and transparency of AI systems, as well as on applying AI techniques to NIST research programs. As part of FARSAIT, NIST's National Cybersecurity Center of Excellence (NCCoE) is examining aspects of trustworthy AI systems and developing a taxonomy that addresses attacks on and defenses for AI systems and technologies. In addition to this kind of general AI research, NIST scientists are applying AI to metrology problems in many domains of national importance, including materials, robotics, and wireless networking.

The National Oceanic and Atmospheric Administration (NOAA) is making long-term investments in AI R&D for many projects related to the oceans and climate. NOAA scientists are applying machine learning to the problem of assessing marine mammal and fish stock abundance from aerial and ship-installed image collection systems and underwater acoustic recordings. Autonomous and semiautonomous robotic systems are being employed for deep sea exploration, including unmanned systems for bathymetric, hydrologic, oceanographic, atmospheric, fishery, ecosystem, and geographic surveys. NOAA scientists also are using machine learning to analyze satellite imagery for predicting oil spill and hazardous material trajectories, wildfire detection and movement, and detection of illegal fishing activity.

NASA applies many technologies, including AI, to problems in its domains of interest: Earth science, aeronautics, human space exploration, and space science. These domains generate significant volumes of data of great variety; NASA uses these data to get as close as possible to reverse-engineering natural systems (e.g., Earth and space) or engineered systems (e.g., aeronautics and human space exploration) to gain insights that result in new machine learning and other tools and technologies. These are then used to help determine future actions to be taken by the systems themselves, in the case of engineered systems, or by human observation systems, in the case of natural systems over which people have minimal control. Many of these actions must be performed with low or no involvement of humans due to the need for quick responses (e.g., real-time operations of unmanned aerial vehicles, or UAVs); the volume, complexity, and/or velocity of data needed to make decisions (e.g., Earth science sensor web control and observation planning, and cybersecurity); low communication bandwidth (e.g., for Mars exploration); and other considerations.

These situations require the use of autonomous systems and motivate NASA's significant push into autonomy, including robotics, planning and scheduling, and machine learning, together with ongoing work in sensor technology and platform development.

Within Earth science, NASA's Advanced Information Systems Technology (AIST) program (see sidebar) funds proposals from teams across NASA, universities, and industry for R&D and implementation of information technologies—including AI—for use in Earth science and downstream applications such as disaster response and water resource monitoring. Within aeronautics, NASA devises new algorithms for anomaly detection with the goal of automatically identifying anomalies within commercial aircraft flight-recorded data and radar track data. Within human space exploration and space science, integrated systems monitoring is needed to allow for health monitoring and maintenance of space systems during operations, due to the lack of human accessibility.

Defense

In 2018, the Department of Defense issued two reports, the *National Defense Strategy* and the *Department of Defense Artificial Intelligence Strategy*,[11] which describe the Department's overall framework for AI development. DoD's goals are to build a more capable military force, strengthen alliances and attract new partners, and reform the Department for greater performance and affordability. The Joint Artificial Intelligence Center established in June 2018 is DoD's focal point for execution of the DoD AI Strategy. For long-term R&D investments, the Office of the Secretary of Defense (OSD)—including the Defense Advanced Research Projects Agency (DARPA)—and the military service laboratories are in the lead.

NASA	***Developments for Earth and Beyond***

NASA's AIST program has led to developments such as deep neural networks for land cover classification (e.g., asphalt, trees, grass, bare ground, or water); assessment of surface deformation due to earthquakes or volcanoes; assessment of coral reefs from multiresolution data through domain adaptation; symbolic regression to learn linear and nonlinear equations describing vegetation as a function of key climate variables; and time series search-and-analysis algorithms, including enhancements of these algorithms through quantum computing.

NASA investments in aeronautics have resulted in a new algorithm for active learning that uses limited domain-expert feedback to reduce the false positive rate of data-driven anomaly-detection algorithms.[9] NASA devised a new algorithm for automatically identifying precursors to known anomalies, where the precursors may happen significantly in advance of the anomalies. A project called Autonomy Teaming and TRAjectories for Complex Trusted Operational Reliability (ATTRACTOR) is developing a large set of technologies to enable explainable AI for multi-UAS (unmanned aircraft systems) missions during mission planning and execution.

NASA's human space exploration and space science activities have yielded clustering and deep learning[10] algorithms for monitoring and assessing the health of key systems such as the International Space Station and the Advanced Multi-Mission Operations System, which is a comprehensive system of tools to be used in future space missions. Automated planning and scheduling have been developed and used extensively for scheduling operations of systems such as the Mars rover, for scheduling crew operations, and for fault management. Machine learning is being developed as part of the project Autonomous Medical Operations for Deep Space Exploration, where the goal is development of medical models that are able to accurately diagnose a variety of pathologies that might occur on deep space missions, derived from ultrasound image analytics and/or machine learning.

[9] S. Das, M. R. Islam, N. K. Jayakodi, J. R. Doppa. Active Anomaly Detection via Ensembles: Insights, Algorithms, and Interpretability. arXiv.org Preprint (2019). https://arxiv.org/abs/1901.08930.

[10] *Clustering* algorithms are used to discover groups of similar items in input data. *Deep learning* is an ML approach that uses simulated artificial neural networks rather than traditional programming or statistical methods. "Deep" refers to use today of large, many-layered artificial neural networks as opposed to the simpler networks studied earlier.

[11] For unclassified summaries, see http://nssarchive.us/wp-content/uploads/2018/01/2018-National-Defense-Strategy-Summary.pdf and https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF.

In September 2018, DARPA announced its "AI Next" campaign, a multiyear investment in new and existing AI R&D programs. Key areas of the campaign include automating critical DoD business processes, such as security clearance vetting or accrediting software systems for operational deployment; improving the robustness and reliability of AI systems; enhancing the security and resiliency of ML and AI technologies; reducing power, data, and performance inefficiencies; and pioneering the next generation of AI algorithms and applications, such as explainability and commonsense reasoning. Currently DARPA is conducting programs applying AI and programs that are advancing the state of the art in AI. DARPA has launched new major programs since the September 2018 kickoff of the AI Next campaign.

AI Next builds on DARPA's five decades of AI technology creation to define and to shape the future. As with all DARPA programs, AI Next programs are designed to create revolutionary advancements in AI and are managed using metrics and independent validation and verification (IV&V) testing to monitor progress and mitigate risk. The AI Next campaign is particularly focused on advancement in the following areas:

- *New Capabilities*: AI technologies are applied routinely to enable DARPA R&D projects, including more than 60 existing programs such as the Electronics Resurgence Initiative and other programs related to real-time analysis of sophisticated cyberattacks, detection of fraudulent imagery, construction of dynamic kill-chains for all-domain warfare, human language technologies, multimodality automatic target recognition, biomedical advances, and control of prosthetic limbs. DARPA will also advance AI and other technologies to enable automation of critical business processes.
- *Robust AI*: AI technologies have demonstrated great value to missions as diverse as space-based imagery analysis, cyberattack warning, supply chain logistics, and analysis of microbiologic systems. At the same time, the failure modes of AI technologies are poorly understood. DARPA is working to address this shortfall with both analytic and empirically focused R&D. DARPA's success is essential for DoD to deploy AI technologies, particularly to the tactical edge, where reliable performance is required.
- *Adversarial AI:* Machine learning systems can be easily duped by changes to inputs that would never fool a human. The data used to train such systems can be corrupted, and the software itself is vulnerable to cyberattacks. These and other areas must be addressed at scale as more AI-enabled systems are operationally deployed.
- *High-Performance AI*: Computer performance increases over the last decade have enabled the success of machine learning, in combination with large datasets and software libraries. More performance at lower electrical power is essential to allow both data center and tactical deployments. DARPA has demonstrated analog processing of AI algorithms with thousand-fold speedup and thousand-fold power efficiency over state-of-the-art digital processors, and it is researching AI-specific hardware designs. DARPA also is attacking the current inefficiency of ML by researching methods to drastically reduce requirements for labeled training data.
- *Next-Generation AI*: This thrust will transform computers from tools into problem-solving partners. DARPA research aims to enable AI systems to explain their actions and to acquire and reason with commonsense knowledge.

In addition to new and existing DARPA research, a key component of the AI Next campaign is DARPA's Artificial Intelligence Exploration (AIE) program, which was first announced in July 2018. AIE constitutes a series of high-risk, high-payoff projects where researchers work to establish the feasibility of new AI concepts within 18 months of award. Leveraging streamlined contracting procedures and funding mechanisms enables these efforts to move from proposal to project kick-off within three months of a funding opportunity announcement. Between July 2018 and October 2019, DARPA has solicited

proposals for 13 different AI topics and has completed making project awards, the majority of which are to universities and small businesses.

AFOSR manages the basic research investment for the United States Air Force (USAF). Within Air Force Research Laboratory (AFRL) activities, technical experts discover, shape, and champion research within AFRL, university, and industry laboratories to ensure that results of Air Force-funded research are transitioned to support USAF needs. AFOSR has two programs dedicated to long-term investment in AI: Computational Cognition and Machine Intelligence (CCMI) and Science of Information, Computation, Learning, and Fusion (SICLF). CCMI supports innovative basic research on the fundamental principles and methodologies needed to enable intelligent machine behavior, especially in support of human-machine teaming. The overall vision of this program is that future computational systems will achieve high levels of performance, adaptation, flexibility, self-repair, and other forms of intelligent behavior in the complex, uncertain, adversarial, and highly dynamic environments the Air Force faces. This program covers the full spectrum of computational and machine intelligence, from cognitively plausible reasoning processes responsible for human performance in complex problem-solving and decision-making tasks, to noncognitive computational models of intelligence necessary to create robust intelligent systems. The SICLF program focuses on the fundamental science needed to "extract and derive needed information" from the large and heterogeneous datasets USAF collects, ranging from radar to video. To understand and interpret information contained in various data sources, it is necessary to extract relevant information from these datasets and make inferences based on prior knowledge and probabilities.

ONR coordinates, executes, and promotes the science and technology programs of the United States Navy and Marine Corps. ONR manages and funds basic and applied science and advanced technology development through grants and contracts with an array of partners in academia, industry, and government in the United States and around the world. ONR's AI-related programs include:

- *Machine Learning, Reasoning, and Intelligence*, which focuses on developing the science base and efficient computational methods for building versatile intelligent agents (cyber and physical) that can perform various tasks with minimal human supervision. In addition, the intelligent agents should be able to collaborate seamlessly with teams of humans and other agents in environments that are unstructured, open, complex, and dynamically changing.
- *Human Interaction with Autonomous Systems*, which aims to create cognitively compatible intelligent autonomous systems and robots that could serve as teammates with humans, enabling peer-to-peer collaboration among humans, robots, intelligent agents, and autonomous systems. Such systems would reduce the cognitive burdens of interfaces with intelligent autonomous systems and enable supervisory control and human training of robots and autonomous systems.
- *Cognitive Science for Human-Machine Teaming*, which seeks to develop empirically grounded models of human cognition and computational architectures that attain human-level intelligence. This research pursues the vision of cognitively and socially compatible intelligent systems that enhance or replace humans and excel as teammates, trainers, and assistants across the spectrum of operational concepts for human-machine teaming.
- *Science of Autonomy*, which focuses on multidisciplinary research topics in autonomy that explore interconnections between fields such as artificial intelligence, machine learning, control theory, human factors, biology, cognitive science, psychology, economics, operations research, applied mathematics, oceanography, physics, and neuroscience.

The Army Research Office (ARO) and its Army Research Laboratory also support R&D to enhance national security applications of AI. In 2018, ARO's work fell into three clusters: (1) autonomous systems, including adaptive sensing, mobility, and navigation; (2) analytics for situational awareness,

including data fusion, artificial neural networks, predictive analytics, and commonsense and adversarial reasoning; and (3) human-AI teaming, including interface technologies, natural language processing, and measurement of soldier training and performance.

Economy

A challenge for U.S. agriculture is producing more food for a growing population on a limited land base while advancing efforts to reduce agriculture's impact on the environment and ecosystems. One of the most promising areas for addressing this challenge is the application of AI in agriculture, food, and related industries, which contribute approximately $1 trillion to the U.S. gross domestic product, roughly a 5 % share. Building on a long history of AI-related research and education activities that span multiple disciplines and program areas, USDA/NIFA is catalyzing efforts that harness the power of AI throughout the food supply chain (see sidebar).

NIFA's data science initiative, the Food and Agriculture Cyberinformatics and Tools (FACT) competitive grants program, focuses on data science to enable systems and communities to effectively utilize data, improve resource management, and integrate rapidly emerging technologies, including AI, to further U.S. food and agriculture enterprises.

AI contributions to the economy include development of autonomous robots, crop and soil sensor data integration into decision support systems, and machine learning models for ecosystems. In recent years, NIFA has invested in competitive grant programs to support AI-related research projects, including:

- Systems for crop and soil monitoring that leverage machine learning, remote sensing, satellite imagery,

NIFA	*Advances in Food and Agriculture*

USDA/NIFA aims to couple better infrastructure for science synthesis with mechanisms for using on-farm data to deliver better recommendations to farmers, e.g., for optimizing fertilizer management with more localized relevance, on-farm value, and compliance with emerging farm-to-fork sustainability metrics. NIFA is developing a cyber-enabled process to link public and private data to create a model ecosystem and community for public-private research and clinical trials for continuous improvement in evidence for management practices. The process will enable best practices, policies, and tools for public-private partnerships for data use in evidence-based practice that addresses privacy, security, cost, and intellectual property concerns; triage and recovery of high-value legacy data; and conversion and sharing of grey/dark data from null or confirmatory studies in FAIR (findable, accessible, interoperable, reusable) formats. Experience gained in conducting data-intensive syntheses of agronomic research and on-farm data will help equip practitioners for site-specific management and use of their data. This NIFA-supported project includes a cyber framework for customization that ingests and anonymizes user data, combines ingested with existing research and on-farm data, and uses AI and ML to fine-tune recommendations with accruing new data to minimize human resource needs.

In animal agriculture, NIFA is investing in the development of a virtual "Dairy Farm Brain" that is intended to be the next big leap in dairy farm management using coordinated data ecosystems. The dairy farming industry is saturated with data; however, data streams are poorly linked to each other and to decision-making processes. NIFA-supported research is addressing this problem using a collaborative, open environment to link data from dairy operations with analytical and other value-added services. This critical work will move the industry toward modernized data exchange and AI-supported decision making.

Retailers reject, discard, or donate some 19.5 million metric tons of edible perishable food products every year—a significant loss of economic, social, and ecological value. NIFA is investing in new research that aims to develop a set of empirical methods applied to data on retail food purchase and sale transactions that can help reduce the amount of food loss and food waste in the retail consumer market. A set of ML algorithms is being developed to assist retailers in better matching demand-flow to wholesale purchases and to analyze online and offline purchase patterns to better understand the implications of online purchasing for household waste. This work advances knowledge of how managing fresh-food supply chains can be more sustainable and profitable for all stakeholders.

drones, and precision technologies for informed production and management.
- Autonomous robots to perform previously labor-intensive tasks like harvesting crops in greater volumes and faster than traditional human laborers.
- Computer-vision algorithms that can distinguish between crops and weeds for robotic weeders, and between fruit and leaves for robotic fruit harvesters.
- Nanotechnology-based sensing mechanisms and smart sensors for accurate, reliable, and cost-effective early and rapid detection of pathogens, allergens, chemicals, and contaminants in foods, plant and animal production systems, water, and soil.
- Development of intelligent decision support systems for farming to increase agricultural production while reducing water use and impact on ecosystems.

In addition to funding AI research projects, NIFA continues to invest in workshops to discuss opportunities and gaps in AI and big data research. NIFA's AI research support is incorporated across several research, education, and extension areas of the agency's flagship competitive grants program, the Agriculture and Food Research Initiative.

Within the Department of Commerce (DOC), the Institute for Telecommunication Sciences of the National Telecommunications and Information Administration (NTIA) is actively researching use of ML and neural networks to perform identification, protection, and classification of wireless signals. This work is a big step towards automating detection of interference events and enforcement of spectrum sharing agreements.

The Department of Transportation's (DOT) Federal Highway Administration (FHWA) supports several large-scale applications of AI to increase the safety, performance, and reliability of the U.S. highway transportation network. The FHWA Safety Data Analysis Team uses the Second Strategic Highway Research Program Naturalistic Driving Study datasets of video and sensor readings to better understand how crashes happen and how people can avoid them. The team is using AI-driven computer vision algorithms to automatically analyze and annotate the video streams with information about work zones, traffic signs, traffic signals, and visibility conditions such as fog or rain. Machine learning algorithms improve the speed and accuracy of labels over those made by human annotators. The Exploratory Advanced Research (EAR) Program is supporting research that classifies large mammals from video images taken by camera traps along highways. Automation will increase the effectiveness of monitoring and could lead to active systems that further reduce wildlife traffic conflicts in real time. The outcomes of earlier EAR projects have been incorporated into tools used by FHWA's Safety Training and Analysis Center.

The Technology Transformation Service of the General Services Agency (GSA) held a special AI-centric funding cycle in 2018. Projects included AI Data Standardization, AI in Civil Service, AI for Website Accessibility, and AI-Assisted Contracts. All of these projects were aimed at making government processes more efficient and more accessible to the general public.

Health

As the primary U.S. Government agency responsible for biomedical and public health research, the Department of Health and Human Services (HHS) supports a large portfolio of AI research. NIH, through projects at 27 institutes and centers, follows a unified research agenda for using AI to turn discovery into health. Several aspects of AI are inspired by neuroscience studies of how the brain computes and processes information. NIH supports extensive research to advance that foundational research via the BRAIN Initiative and other programs.

HHS AI projects include using AI systems for screening regulated products entering the United States; summarizing drug-labeling documents using deep-learning natural language processing; detecting diseases using medical images processed through novel deep learning architectures (e.g., detecting age-related eye diseases; see sidebar); predicting antimicrobial resistance from whole genome sequencing data; and speeding drug development by AI-based trial simulation and dose optimization. NIH also uses AI approaches to guide evaluation of the over 80,000 research proposals it receives annually.

NIH	*Detecting Age-related Eye Disease*

NIH Researchers at the National Library of Medicine (NLM) and National Eye Institute (NEI) have developed an AI algorithm, called DeepSeeNet,[12] to assist in the screening of age-related macular degeneration (AMD): a leading cause of vision loss in Americans aged 60 years and older. Typically, in assessing the severity of AMD, NEI's Age-Related Eye Disease Study (AREDS) Simplified Severity Scale predicts the risk of progression to late AMD. However, its manual use requires the time-consuming participation of expert practitioners. By leveraging cutting-edge deep learning techniques and repurposing "big" imaging data from a major AMD clinical trial, NIH researchers developed a novel data-driven approach for automated AMD diagnosis with its performance exceeding human ophthalmologists (retinal specialists in this case), highlighting the potential of deep learning systems to assist early disease detection and enhance clinical decision-making processes.

Additionally, at the NIH National Library of Medicine (NLM), which conducts and supports research as well as provides access to a vast array of literature and scientific data banks, AI and machine learning methods provide fast and efficient access to the literature, delivering relevant results[13] to user queries within seconds. These methods are used in NLM's PubMed repository, comprised of 30 million biomedical literature citations, searched daily by millions of people. Further, the advisory committee to the NIH Director and an AI working group continue to provide guidance in identifying opportunities for cross-NIH efforts that reach broadly across biomedical topics, building bridges between the computer and data science communities and the biomedical research community, and developing ethical guidelines for use of AI in health research and healthcare.[14]

In addition to the technical work of NIH, the Food and Drug Administration (FDA) of HHS supports regulatory science efforts to ensure the safety and effectiveness of medical devices and products that rely on AI. FDA is also developing appropriate analytical and clinical testing methods, standards, and performance benchmarks and comparators for appraising the safety and effectiveness of medical devices and products.

Justice & Security

The Department of Homeland Security (DHS) is exploring a wide range of emerging AI technologies and capabilities to advance critical homeland security missions. Application areas include cybersecurity, counterterrorism, border security, disaster response, and protection of critical infrastructure.

[12] Y. Peng *et al*. DeepSeeNet: A deep learning model for automated classification of patient-based age-related macular degeneration severity from color fundus photographs. *Ophthalmology* (2018) S0161-6420(18), 32185-7. PMID: 30471319.

[13] N. Fiorini, *et al.* Best Match: New relevance search for PubMed. *PLoS Biol*. (2018) 16(8):e2005343. doi:10.1371/journal.pbio.2005343. PMID: 30153250.

[14] See https://acd.od.nih.gov/documents/presentations/06132019AI.pdf and https://acd.od.nih.gov/meetings.html.

Among DHS priorities in support of AI research are anomaly detection; extracting meaning from diverse information sources including video, audio, text, and databases; and triaging data in real or near real-time into actionable intelligence in a variety of different circumstances (e.g., see the sidebar). The DHS Science and Technology (S&T) Directorate is examining technical areas that are the focus of enduring research activities, including advances in real-time analytics and computer vision, as well as supporting policy research to address automation and autonomy in high-consequence decision making.

The Intelligence Advanced Research Projects Activity (IARPA) addresses some of the most difficult challenges in the Intelligence Community by sponsoring research programs, studies, and challenges. Recent and current IARPA research includes about 80 different efforts that use or improve AI. Some of its largest efforts are:

- *Creation of Operationally Realistic 3D Environment (CORE3D)*, which uses machine learning to construct a fully high-fidelity 3D model of the world using remote sensor data.
- *Deep Intermodal Video Analytics (DIVA)*, which uses machine learning to watch streaming video from multiple cameras and automatically detect when people and vehicles take specific actions.
- *Finding Engineering-Linked Indicators (FELIX)*, which uses machine learning to learn signatures of genetic engineering across multiple organisms.
- *Functional Genomic and Computational Assessment of Threats (Fun GCAT)*, which uses machine learning to classify genetic (e.g., DNA) sequence data by genetic taxonomy, sequence function, and threat potential.

DHS	***Using AI to Aid First Responders***

Imagine a first responder answering the call to a natural disaster, a house fire, or an active shooter incident where there may be multiple injuries and unknown casualties. The sheer volume of data from command and dispatch centers, cameras and other sensors, and people at the scene risks overwhelming or distracting first responders from their critical activities. As a result, first responders may make the wrong decisions, not because they lack relevant data, but because they are unable to quickly extract key insight from this flood of information. Information overload creates obstacles for first responders to perform their duties safely and efficiently.

In response to this challenge, DHS's S&T Directorate in collaboration with NASA is developing AUDREY, a state-of-the-art human-like reasoning system designed to assist first responders in synthesizing high-level data at the scene of an emergency.[15] It provides situationally relevant information and advanced analytics insight to decision makers and first responders, enabling them to make the best decisions possible in situations where time is critical. AUDREY does not take the place of a human in the decision chain; rather, it is a digital assistant that helps people make better decisions.

Similar to the voice-activated device on a smartphone or a voice-controlled intelligent personal assistant, AUDREY is personalized to the individual responder and can recognize first-responder specialized language. Enabling easy hands-free communication is just the simplest of its functions. The system has knowledge of emergency procedures and tracks the entire team of emergency responders and can proactively send relevant alerts to individuals to help them work together. For example, during a fire, it can help ensure that the team explores a building for survivors quickly and efficiently. AUDREY also integrates information from wearable sensors that monitor a responder's health, such as oxygen and carbon monoxide levels and respiration rates. It can warn responders who are nearing their limits of physical endurance, and when necessary, dispatch help from other responders. New operational processes can be input to AUDREY through a knowledge editing tool. A pilot deployment of the system began in June 2019.

[15] https://www.dhs.gov/publication/st-frg-audrey

- *Machine Intelligence from Cortical Networks (MICrONS)*, which improves machine learning capabilities by reverse-engineering the brain's algorithms, accomplished by neuroimaging studies unprecedented in detail and scale.
- *Multimodal Objective Sensing to Assess Individuals with Context (MOSAIC)*, which is estimating and predicting individuals' psychological, cognitive, and physiological constructs and overall job performance using data from various individual, environmental, and social sensing data streams.
- *Rapid Analysis of Various Emerging Nanoelectronics (RAVEN)*, which uses machine learning to accelerate the speed and accuracy of image processing for state-of-the-art integrated circuits.

Within the Federal Bureau of Investigation (FBI), AI systems support the work of many of the Bureau's divisions, including criminal and intelligence missions, and FBI usage of AI systems will continue to increase as the FBI deals with ever-larger datasets and advanced technological capabilities. The FBI funds multiple developments with third-party vendors that incorporate machine learning, deep neural networks, and augmented intelligence in tools to further its mission. This work ranges from investigative and intelligence information to defensive situations that seek to mitigate risk. The Operational Technology Division's Technology Assessment and Advisory Group focuses on developing advanced analytic techniques to manage the volumes of data collected for investigations. The FBI also supports R&D through academic fellowship programs and direct engagement with industry.

Strategy 2: Develop effective methods for human-AI collaboration

Many Federal agencies are undertaking efforts to foster the development of effective methods for human-AI collaboration. While these efforts align with specific organizational missions, they also cluster around generalizable topics such as human-machine cognition, autonomy, and agency in the contexts of trusted machine intelligence, decision support, situational awareness, and risk modeling.

Science & Engineering

NSF has many programs that address human-AI interaction and collaboration. These include core programmatic areas in NSF's CISE and SBE Directorates, as well as a growing number of interdisciplinary programs. The CISE Cyber-Human Systems core program funds research on AI systems that augment human physical and mental capabilities, enable more fluid human-machine interaction by sensing the human's body language and emotional state, and enhance problem-solving with mixed teams of people and AI systems. The Behavioral and Cognitive Science Division of the SBE Directorate supports research on human perception, decision making, and use of language. The models of human cognition created by this research are influential in the design of AI systems that interact with humans.

The interdisciplinary, interagency NRI 2.0 program supports fundamental research that accelerates development and use of co-robots that work cooperatively with people. Exemplary research thrusts in NRI 2.0 include:

- *Scalability*: How robots can collaborate effectively with multiple humans or other robots; how robots can perceive, plan, act, and learn in uncertain, real-world environments, especially in a distributed fashion; and how to facilitate large-scale, safe, robust, and reliable operation of robots in complex environments.
- *Customizability*: How to enable co-robots to adapt to specific tasks, environments, or people with minimal modification to hardware and software; how robots can personalize their interactions with people; and how robots can communicate naturally with humans both verbally and nonverbally.
- *Lowering barriers to entry in robotics*: How to lower the barriers to conducting fundamental robotics research and research on integrated robotics applications. This may include development of open-source co-robot hardware and software, as well as widely accessible testbeds.
- *Societal impact*: How to establish and infuse robotics into educational curricula; to advance the robotics workforce through education pathways; and to explore the economic, ethical, legal, and societal implications of a future with ubiquitous collaborative robots.

NIST is researching human-AI collaborations and interactions in the area of usability at both fundamental and applied levels. Usability is defined as efficiency, effectiveness, and user satisfaction in a specific context of use. NIST's usability research will fill a gap where there are currently no established methods for quantitatively measuring the quality of human-AI interactions. Human-AI collaboration research is also taking place in the application area of manufacturing robotics.

In addition, NIST is developing means of quantifying the impact of robot appearance and behavior on the effectiveness of human-robot relations and interactions in order to drive future efforts in metrology for teaming-related AI and artifact designs.

NOAA is advancing human-AI collaboration for a wide variety of tasks, including:

- Combining traditionally manual marine mammal and fish stock abundance surveys with aerial and ship-installed image collection systems and underwater acoustic recordings.
- Employing remotely operated robotics with automated control of routine dive sequences for deep sea exploration.

- Assisting analysts with machine learning to use satellite imagery for severe weather detection and prediction, oil spill and hazardous material trajectories, wildfire detection and movement, ecosystem health, and detection of illegal fishing activity.

DOE is the steward for several large scientific user facilities in the National Laboratory complex. In order to make optimal use of these facilities, DOE is employing AI and automatic decision-making methods to automate steering of large experiments. For example, the Center for Advanced Mathematics for Energy Research Applications (CAMERA) has developed a ML system that reduces the number of X-ray crystallography trials needed to determine the shape of a molecule by a factor of six.[16]

Defense

DARPA envisions a future in which machines are more than just tools that execute human-programmed rules or generalize from human-curated datasets; rather, machines will become partners in problem solving (e.g., see sidebar). Incorporating these technologies into military systems that collaborate with warfighters will facilitate better decisions in complex, time-critical, battlefield environments; enable shared understanding of massive, incomplete, and contradictory information; and enable unmanned systems to collaborate with human warfighters to perform missions more safely and with higher degrees of autonomy.

AFOSR's Trust and Influence Program is dedicated to long-term investment in human-machine teaming. The program's first goal is to advance basic understanding of human reliance and teaming to elucidate how people establish, maintain, and repair trust in agents, both human and machine. In particular, the program supports research to build the scientific foundations for

DARPA	***XAI: Explainable Artificial Intelligence***

AI systems learn to solve problems such as classifying inputs (e.g., is this a picture of a cat or dog?) or making decisions (e.g., what treatment should be given to a patient with particular symptoms?) by automatically generalizing from a large set of examples. The inner workings of such a system are a "black box": given a question, the system provides an answer, but it cannot explain how or why it reached that answer. Lacking the ability to explain itself, the system cannot be trusted in high-stakes applications, particularly because machine learning systems occasionally make catastrophic mistakes. The DARPA Explainable Artificial Intelligence (XAI) program aims to improve trust and collaboration between humans and AI systems by developing ways for the systems to explain their reasoning in human-understandable terms. In addition to enhancing trust, explanations also allow people to spot and correct errors the AI system made when it generalized from its training data. In the envisioned virtuous cycle of AI-human collaboration, the AI system helps humans solve problems more quickly, and human oversight and feedback helps the AI system become "smarter" and more accurate.

XAI has supported R&D on a wide range of approaches. Some of the most exciting and significant results have shown how an artificial neural network image classification system—historically one of the blackest of black boxes—could be designed to explain its answers in plain English by referring to features in the image. For example, when the system classified a picture of a bird as that of a downy woodpecker, it stated, "This is a downy woodpecker because it is a black and white bird with a red spot on its crown." Most importantly, the systems were not canned responses but were generated on the fly by another artificial neural network that was tied to and co-trained with the image classifier. The same system can be used for critical tasks ranging from classifying diseases from medical images to classifying combatants and civilians in military surveillance images. The work is an important step toward helping AI systems and people work together seamlessly and confidently.

[16] "Autonomous Steering of X-Ray Scattering Experiments Through Optimization and Artificial Intelligence," https://www.camera.lbl.gov/copy-of-single-particle-imaging-for.

designing high-performing, mixed human-machine teams through properly calibrated trust. Of specific interest is research on the human social and cognitive processes that inform the design of systems composed of human and machine agents and the development of novel man-machine interfaces and interaction techniques. Another program goal is to advance the science of social influence within the context of national security. For instance, researchers in the program strive to understand the variables that influence human behavior, attitudes, and beliefs sufficiently well for measurement and forecasting of social phenomena. There is particular interest in developing computational approaches and in using large-scale datasets to understand social and cultural behavior. The Trust and Influence Program invests in the discovery of foundational concepts of effective influence, deterrence, trust-building, trust calibration, and counterterrorism operations.

Economy

FHWA is investigating the use of cognitive agents to assist human travel services. AI-based tools involving cognitive agents are being used to enhance the human services side of transportation by making inferences about specific travel and mobility needs in transportation facilities. This is especially important for information sharing, such as for passengers with disabilities traveling across multiple modes of transportation. AI technologies also can aid in decision making for traffic management through incident detection and demand forecasting. Expert systems will increase knowledge of individual localities and assist in coordination across jurisdictions.

Health

A major focus of NIH work on human-AI collaboration is on document processing for knowledge management. A significant amount of documentation and data resides within the agency in the format of word-processing documents. NIH researchers are developing systems that can extract information from massive numbers of free-form natural language documents, recognize and translate foreign languages on the fly, and summarize key information in tabular form. The resulting tables can be quickly read by doctors and scientists or used as inputs for further analysis.

Toward this goal, the NLM has issued a notice of special interest[17] for soliciting grant applications through NLM Research Grants in Biomedical Informatics and Data Science for state-of-the-art methods and approaches to address problems with large health datasets or tools used to analyze them. These data could be drawn from electronic health records or public health datasets, biomedical imaging, omics repositories, or other biomedical or social/behavioral datasets. In addition to addressing technical advances, applicants were encouraged to address ethical issues that might arise from their proposed approaches.

DOE has a long history of supporting research on human-AI collaboration to improve environmental health issues resulting from nuclear waste storage and clean up. DOE and its liquid nuclear waste contractors routinely develop and deploy advanced semiautonomous robotic systems for tasks such as monitoring radiation levels, retrieving samples, and repairing leaks.[18] In addition to cleaning up after problems, DOE researchers have also begun to explore the use of machine learning to help provide early warnings of potential problems such as groundwater contamination.[19]

---

[17] https://grants.nih.gov/grants/guide/notice-files/NOT-LM-19-003.html

[18] SRS Relies on Array of Robots for Liquid-Waste Cleanup, DOE Department of Environmental Management, Aug. 31, 2017, https://www.energy.gov/em/articles/srs-relies-array-robots-liquid-waste-cleanup.

[19] F. Schmidt, H.M. Wainwright, B. Faybishenko, M. Denham, C.l. Eddy-Dilek. In Situ Monitoring of Groundwater Contamination Using the Kalman Filter. *Environmental Science & Technology* (2018) 52(13), 7418-7425.

Justice & Security

The National Institute for Justice (NIJ) is funding a number of R&D projects to apply advances in AI to reduce crime and advance justice. The Cognitive Agent project aims to improve community supervision of people on parole by automatically detecting potentially risky behaviors from time-stamped location data and alerting both the subjects and their parole officers. Another project is developing a mobile system to provide offenders with increased access to personalized real-time resources and opportunities to help them successfully transition to life outside prison. AI technology for planning and scheduling is being used to balance officer caseloads. Finally, NIJ is exploring ways to use AI technology to automatically search for and analyze investigative leads, in particular for Internet-based crime.

DHS is addressing human-AI collaboration across several operational programs. It uses AI to fuse situationally relevant information in real time and provide advanced analytics insight to decision makers and first responders. AI also is engaged in improving emergency preparedness and increasing responder effectiveness in active disasters by improving decision making, reducing response time, and mitigating disaster consequences. Additionally, DHS engages in AI R&D to improve quality of detection and other information, including enhancing capabilities to quickly triage diverse data sources, to provide real-time or near real-time situational awareness, predictive scoring, and entity resolution.

At IARPA, the Machine Translation for English Retrieval of Information in Any Language (MATERIAL) program identifies foreign-language information from speech and text relevant to English queries and provides meaningful evidence of relevance of the retrieved information in English. A related program, the Open Cross-Language Information Retrieval (OpenCLIR) Challenge develops ML methods in a low-training data condition to retrieve Swahili speech and text documents relevant to English queries.

Strategy 3: Understand and address the ethical, legal, and societal implications of AI

The U.S. Government has a number of research and development initiatives that will increase understanding of the potential ethical, legal, and societal implications of AI technologies, which will inform discussions of how they are best addressed. Issues that U.S. Government R&D funding is being used toward include fairness, transparency, privacy, reliability, resilience, algorithmic bias, accountability, and explainability. The information gathered in this research can be used to better understand how AI systems work, potential challenges, and how those challenges will evolve as the technology progresses, how the challenges can be mitigated, and the broader societal impacts of these revolutionary technologies and design solutions. Current research is drawing on law, social science, and ethics (including privacy, civil rights, and civil liberties), as well as expertise in specific technologies and application domains. U.S. Government research is being done in the following categories:

- *Design, development, deployment, and use:* Examines assumptions, appropriate limitations on use, and security controls for sensitive personal data with AI use.
- *Monitoring and testing:* Includes continual assessment and validation of inputs and outputs as to their continued compliance with laws, avoidance of bias, transparency, and explainability.

A number of Federal agencies have developed R&D programs to explore the ethical, legal, and societal implications of AI. These agencies are working to build expertise, bridge interdisciplinary gaps, and form relevant research agendas to tackle these challenges.

NSF	***Making Fair Predictions***

AI systems are typically trained to maximize the accuracy of their predictions, but they may then exhibit discriminatory behavior based on sensitive characteristics such as race and gender. AI researchers have tried to combat this problem by building into the AI system a requirement that it be fair, for some quantitative definition of fairness. Fair AI implementations typically have been assessed by how they behave in one-off prediction tasks, without regard for how the predictions change affected populations of people over time.

In 2018, NSF-funded researchers studied the long-term impact of AI systems on different groups in a population when the AI made predictions that affected the population itself. The study used a one-step feedback model to examine how static fairness criteria interact with temporal indicators of well-being, notably long-term improvement, stagnation, or decline. The researchers showed that such fairness criteria in general do not promote improvement in well-being over time and may in fact cause harm in certain cases.[20] These theoretical results were corroborated by experiments on FICO® credit score data from 2003, which showed that under various models of bank utility and score change, the outcomes of applying static fairness criteria would occasionally make particular decisions less fair. This study greatly elevated the quality of analysis of the full ethical, legal, and societal implications of AI.

Ethical & Legal Considerations

Several NSF programs (e.g., see sidebar) have called for projects focused on the fairness and ethical implications of AI and techniques to mitigate bias and enhance accountability, transparency, and robustness to ensure societal benefit.[21,22,23] NSF and a major American technology company are collaborating to support

[20] L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, M. Hardt. Delayed Impact of Fair Machine Learning. *Proceedings of Machine Learning Research* (2018) Vol. 80, 3150-3158.

[21] https://www.nsf.gov/pubs/2019/nsf19018/nsf19018.jsp

[22] https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505651

[23] https://www.nsf.gov/pubs/2019/nsf19016/nsf19016.jsp

research focused on AI fairness with the goal of contributing to AI systems that are readily accepted and deployed to tackle grand challenges facing society.[24] Specific topics of interest include, but are not limited to, transparency, explainability, accountability, potential adverse biases and effects, mitigation strategies, validation of fairness, and considerations of inclusivity.

Use of AI for law enforcement and national security is raising complex ethical and legal questions. To keep ahead of the curve and ensure that use of AI does not unfairly or illegally disadvantage individuals or groups, DHS is applying and extending its existing frameworks and tools, including the 2012 *Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*,[25] the 2013 companion report *Applying Ethical Principles to Information and Communication Technology Research*,[26] and the online Cyber Research Ethics Decision Support tool.[27] DHS's S&T Data Analytics Technology Center (DATC) has identified bias and fairness in AI systems as priority issues and supports both qualitative and quantitative research in these issues.

DoD's Artificial Intelligence Strategy also calls out leadership in military ethics. DoD is investing in the research and development of AI systems that are resilient, robust, reliable, and secure. It funds research into techniques that produce more explainable AI; and it is pioneering approaches for AI test, evaluation, verification, and validation. It will develop and implement AI principles for defense in formal policy. DARPA, for example is funding work that explores the cultural implications of bias in datasets and has embedded an Ethical, Legal, Social Implications panel in the technical development of future autonomous systems that may be deployed in complex urban environments.

DOE is working on the emergent privacy, ethics, and adjacent impacts that may arise in applications of AI to ever larger aggregations of data and information. There is an initial focus based on network traffic and data related to cybersecurity.

Societal Implications

AI will have important impacts on the economy, the workforce, and the nature of work. Several Federal agencies have begun to study these impacts. For example, in 2018, NSF awarded a first round of grants under *The Future of Work at the Human-Technology Frontier* (FW-HTF) Big Idea program[28] to understand the impacts of AI in a range of workplace settings.

NIFA is funding research into the social implications of AI for farm, forest, and ranch managers and is studying how AI could help create economic opportunities for America's rural communities and entrepreneurs.

The deployment of AI across society could also affect homeland security in complex and unpredictable ways. DHS DATC is funding interdisciplinary research on homeland security concerns related to AI, such as the potential use of AI systems by malicious actors, AI's potential to change communities in ways that impact homeland security missions, and potential challenges related to risk perception and communication.

DOT is exploring how AI could help promote road safety. The FHWA EAR program is helping develop AI systems that can study huge troves of video-based safety data while preserving driver privacy through

[24] https://www.nsf.gov/pubs/2019/nsf19571/nsf19571.htm
[25] https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf
[26] https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCOMPANION-20120103-r731_0.pdf
[27] https://dl.acm.org/citation.cfm?doid=2793013.2793017
[28] https://www.nsf.gov/eng/futureofwork.jsp

automatic facial masking. Other AI tools developed through the EAR program are refining automated feature extraction of video to identify roadside features, the presence of other motorized and nonmotorized roadway users, and the relative distance of all objects from the camera's point of view. These AI-assisted approaches may unlock opportunities for new types of road safety research, which could prove especially valuable as America's transportation system becomes increasingly automated.

Strategy 4: Ensure the safety and security of AI systems

AI safety encompasses securing AI systems as well as the digital infrastructure upon which these systems run. Federal cybersecurity efforts are more advanced than AI safety efforts. Numerous Federal agencies are beginning to use AI to strengthen the cybersecurity of systems, including AI systems, and other key digital assets. Nevertheless, Federal agencies have taken promising early steps toward separately assuring the safety of AI systems. A robust new set of Federal AI safety programs will ensure that the United States is in the strongest possible position to safely and securely unlock the full potential of AI (see sidebar).

AI Safety

Many cutting-edge AI systems, especially systems based on machine learning methods, are difficult or impossible to understand in action—the system may be able to complete a specific task, but the internal "decision criteria" it uses to do this often cannot be isolated or explained in terms a human can understand. This black box problem makes it difficult, even impossible, to pinpoint and fix vulnerabilities in many of the most useful AI systems available today, making AI "explainability" a key AI safety challenge. Federal agencies are working to tackle this challenge. In January 2019, NSF and several leading universities established the Center for Trustworthy Machine Learning, a five-year effort to examine security issues such as these.

As noted in a sidebar in the Strategy 2 section, DARPA's XAI program is exploring technical means to make AI systems more understandable for humans while still ensuring high performance. XAI aims to better understand AI vulnerabilities and foster greater user confidence in AI systems.

In addition, DARPA's Guaranteeing AI Robustness Against Deception program is researching how to build AI systems that are resistant to adversaries' efforts to fool them—for example, by deploying optical illusions against an AI image recognition system to cause the system to malfunction.

NSF	***Protecting AI Against Adversaries***

Modern AI systems are often built using deep neural networks, which are susceptible to AI-specific camouflage called "adversarial examples". Since adversarial examples were discovered several years ago, there have been increasing numbers of efforts to construct defenses to combat adversarial networks and keep neural networks robust against attackers. This is particularly difficult in a "white box" threat situation, where the attacker has full access to the model, including the neural network architecture and weights; the training algorithm and training data; and any randomness introduced into the model when making predictions. In 2018, nine different studies proposing defenses to adversarial examples were published at a top machine learning conference. NSF-funded researchers built new techniques to defeat seven of the nine published defenses in the few months between when the studies were published and the conference occurred.

The researchers identified a shared pattern across the proposed techniques for defending against adversarial examples: "obfuscated gradients". The defenders were trying to keep attackers at bay by obfuscating how an AI system changed its responses as a result of a change in inputs. This change or gradient is normally the information that an attacker uses to search to find an input that is an adversarial example against the AI system. By obfuscating the gradient, the attacker would be unable to find an adversarial example. While this is an intuitive defense, NSF-funded researchers showed that it is insufficient, as they were able to recreate the gradients and use them to find adversarial examples. They illustrated the depths of the problem of adversarial examples and the ease with which apparently strong defenses can be defeated.[29] Their efforts made clear that AI security will be a long-lasting problem that will require persistent, detailed analyses to address.

[29] A. Athalye, N. Carlini, D. Wagner. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. *Proceedings of the Thirty-Fifth International Conference on Machine Learning (ICML)*, Stockholm, 2018.

IARPA's Trojans in Artificial Intelligence (TrojAI) program is funding projects to detect if an adversary has inserted a Trojan or backdoor into an AI system, such by "poisoning" the training data used for machine learning.

Cybersecurity and AI

Over the past several years, cybersecurity has become a priority in nearly every Federal agency (see examples in sidebars). Researchers are pushing the bounds of cybersecurity by using AI systems to help predict, identify, and deter cyberattacks, including attacks on other AI systems. NSF's Secure and Trustworthy Cyberspace program, IARPA's Virtuous User Environment program, and ongoing NIST and NASA initiatives all aim to use AI to improve detection of cyber anomalies. In 2016, IARPA launched its Cyber-attack Automated Unconventional Sensor Environment program, which seeks to develop novel, automated methods for detecting and forecasting cyberattacks appreciably earlier than existing approaches.

These AI-driven efforts to predict and detect attacks are supplemented by projects that use AI to assist with cyberdefense. DARPA, IARPA, and NSF all have funded multiple projects involving use of AI for cyber defense. NIST researchers are investigating how AI systems can be used to automate vulnerability assessments for digital infrastructure and to produce vulnerability ratings using the industry-standard Common Vulnerability Scoring System.

DOE coordinates R&D projects that develop AI capabilities to enhance the cybersecurity of critical energy systems. These R&D projects include a requirement to verify and validate emerging AI tools and technologies with an energy sector asset owner/operator demonstration partner. To avoid the risk of

DOE	***Protecting AI Against Adversaries***

DOE has been funding research on AI to improve cybersecurity to help keep AI systems secure. Much of cybersecurity today relies on identifying previously seen attack types by detecting hand-coded signatures. A better system, called anomaly detection, is to identify previously unseen attacks by detecting if anything is unusual. Until recently, anomaly detection approaches suffered from prohibitively high false alarm rates and typically leveraged only a small fraction of available data, thereby limiting their detection capability.

DOE-funded researchers addressed this problem by using large quantities of sensor data (real and simulated) and deep reinforcement learning to train an AI system that could detect novel threats with few false alarms by adapting to new cyber threats. The researchers developed a limited proof-of-concept demonstration of this approach[30] and showed that it can (1) autonomously and adaptively choose beneficial cyber defense actions to take, given real-time sensor measurements, and (2) rapidly detect threats with few false alarms by optimally performing these actions over time in response to new real-time data. The system accepted as input raw host-based, sensor-like data. It was then trained on an agent-based simulation of host activity that was fit to historical host-based sensor data and included a memory mechanism for detecting changes in behavior. The researchers demonstrated the benefits of the approach using exfiltration as a use case in a setting where key decisions include if/when to collect additional data from disparate sources. Their efforts show how using AI for cyber defense is a promising pathway to securing many digital infrastructure parts, including AI systems.

NIST	***Understanding AI Attack Models***

Essential to securing AI systems is detailed understanding of the attack models and their countermeasures. NIST's Information Technology Laboratory and National Cybersecurity Center of Excellence have been conducting research to develop a taxonomy of concepts and terminology for adversarial ML that categorizes adversarial capabilities and illustrates potential attacks in the lifecycle of an ML system (e.g., data poisoning during ML training or evasion attack during ML testing). The taxonomy and terminology are intended to inform future standards and best practices for assessing and managing vulnerabilities of ML technologies.

[30] This as-yet unpublished deep reinforcement research at Lawrence Livermore National Labs is described broadly at https://www.llnl.gov/news/playing-video-games-may-help-researchers-find-personalized-medical-treatment-sepsis.

inadvertent disruption in energy delivery, asset owners/operators provide access to testbeds that accurately replicate their operational infrastructure, allowing for validation of the developed technology without compromising safety and reliability.

Strategy 5: Develop shared public datasets and environments for AI training and testing

Government data are as massive, broad, and diverse as the government itself. Data useful for training and testing AI systems include both data collected for programmatic purposes and data created through research; these data come in both structured (database) and unstructured (natural language text) form. Data describe natural phenomena and human activities at a wide range of scales, from the molecular to the astronomical. Agencies are striving to exploit these data for improving the performance of government programs; for scientific discovery, for example, by academics; and for use by the private sector to directly strengthen the U.S. economy. Regardless of data characteristics and AI applications, Federal agencies are facing similar paths to create new approaches for data curation that may include data acquisition, preparation, annotation, integration, and preservation. Beyond the data, agencies are working with nongovernmental partners to create environments necessary to support diverse and large datasets, design new ontologies that make data easier to discover, and provide scalable solutions for data curation.

Science & Engineering

Since 2011, NSF has required all proposals to include plans for data management and sharing of the products of research. Each proposal to NSF must include a data management plan describing how the applicant will conform to NSF policy on the dissemination and sharing of research results, as specified in NSF's *Proposal and Award Policies and Procedures Guide*. The proposal may include the types of data, samples, physical collections, software, curriculum materials, and other materials to be produced in the course of the project; the standards to be used for data and metadata format and content; policies for access and sharing, including provisions for appropriate protection of privacy, confidentiality, security, intellectual property, or other rights or requirements; policies and provisions for reuse, redistribution, and the production of derivatives; and plans for archiving data, samples, and other research products, and for preservation of access to them.

DOE maintains extensive stores of multivariate data to drive AI-based solutions across the economy. DOE platforms provide both secure and public access to data libraries for developing, training, and testing AI-based tools. This includes innovative, secure, public-private virtual environments for multi-entity teams to develop discoveries using proprietary or sensitive datasets—prior to publishing useful results. AI analyses resulting from these to date include accelerating the creation of high-performing catalysts; improved electric grid operations; real-time visualization and analysis of the subsurface for fluid disposal and domestic resource discoveries; and discovery of new particle-physics, drugs, and materials breakthroughs.

GSA hosts the largest publicly accessible database of government data. Data.gov contains over 246,000 datasets from across agencies. Designed to automatically harvest open datasets released by agencies, the repository is widely used by researchers for training and testing AI systems.

All of NASA's systems and associated sensors, across its full range of application domains, generate data. NASA and its partners, collaborators, and the public expect to use these data to yield knowledge on how these systems work. To the extent possible, NASA makes these data publicly available to facilitate non-NASA efforts to derive knowledge from these data. For example, NASA's Earth science data are accessible largely through Distributed Active Archive Centers (DAACs).

The NASA Earth eXchange (NEX-AI) (see sidebar) also has significant Earth science data, as well as associated supercomputing assets, code, publications, and other resources. NASA's AIST program funds combinations of NASA and non-NASA teams to develop information technologies that facilitate use of Earth science data to advance Earth science knowledge. Within space science, data from the Kepler mission and the Transiting Exoplanet Survey Satellite to identify exoplanets, and other missions, are available.[31] NASA's GeneLab database contains human space exploration-related "omics" data and shares those datasets and the results of their analysis with scientists worldwide. In aviation, NASA maintains the Aviation Safety Reporting System, where redacted versions of reports associated with safety events in U.S. commercial aviation are available. These reports have been used by NASA, the Federal Aviation Administration, and others to identify classes of problems that have appeared in commercial aviation. DASHlink is a NASA platform to facilitate collaboration among researchers in data mining and system health management—especially in aviation. More broadly, the NASA Big Data Working Group is concerned with addressing the problem of making datasets from NASA projects available and easily usable after the projects have ended, through standardization of data formats, documentation, tools, results of analyses, and other resources related to these datasets.

NASA	***NEX-AI: NASA Earth eXchange***

The NASA Earth eXchange (NEX) is a large public database of Earth Science data, code, and publications. NEX-AI extends NEX with an end-to-end framework for performing inferences and predictions from Earth Science data using a benchmarked suite of deep learning models, unique labeled training datasets, and workflow automation for both high-performance computing centers and commercial cloud computing services. It goes beyond traditional data repositories by providing a complete work environment—a concept called "science as a service".

NEX-AI provides a powerful and flexible deep learning system named DeepSAT for analyzing earth images from satellites and aircraft. DeepSAT can segment and classify an earth image into different regions based on the kind of ground cover, such as forest, agricultural fields, roads, water, and so on, at a fine level of detail. For example, the code can be used to automatically label regions in satellite images of earth where deforestation is occurring and measure the progression of deforestation over time. The system continually learns and improves over time and can leverage raw unlabeled data—a mode called "unsupervised learning"—as well as data that has been labeled by experts. In addition to DeepSAT, NEX-AI contains SATNet, a database of 5 million labeled images together with a "model zoo" of pretrained deep learning models for different satellites. Researchers can easily modify the provided algorithms or run completely new code and perform detailed and accurate comparisons between algorithms. The NASA NEX-AI team is currently collaborating with leading industry experts in testing newer AI algorithms and in defining hard problems in the land-climate-atmosphere continuum that can possibly be solved by combining multiple kinds of machine learning algorithms.

NIST has created shared datasets and environments spanning robotics, material discovery, wireless spectrum analysis, and semantic annotation. For example, in robotics, NIST is developing methods for generation and curation of AI training data scoped for selected key technical aspects of manufacturing robots. In advanced materials discovery, NIST is creating a high-fidelity database, Joint Automated Repository for Various Integrated Simulations, density functional theory (JARVIS-DFT), with more than 30,000 materials and 500,000 properties to be used as training data. In wireless spectrum analysis, NIST's Spectrum Sharing Models and Measurement Tools project will create a curated radio frequency (RF) signal database to aid in the development of machine learning models for signal detection and classification. Example datasets include radar signals like those planned for the 3.5 GHz band, comprise

[31] https://exoplanetarchive.ipac.caltech.edu/

noise and interference, and can be used to train and evaluate AI detectors enabling Federal-commercial spectrum sharing. In semantic data annotation, NIST is developing and deploying sophisticated AI tools. Semantically labeled datasets are essential for proper support of machine learning and other AI algorithms, particularly when combining data from diverse sources. NIST has also developed and deployed a data storage and dissemination platform, the NIST Public Data Repository.

In addition to datasets, NIST has deployed a new environment for AI training and testing. In 2018, it acquired a mid-scale computing facility optimized for big data and machine learning applications. Of particular interest at NIST is the use of machine learning for the design and optimization of experimental systems in real time.

Because NOAA's mission involves providing data and services, NOAA has developed shared datasets and environments for most of its AI applications. Examples include the Video and Image Analytics for the Marine Environment open source toolkit used for integrating and analyzing aerial and underwater surveys for fisheries and marine mammals; CoralNet, an automated coral classification tool for data collected from benthic surveys; EcoCast, which automates estimation of fishing success for a geographic area and time using satellite and historical data; Automated Surface Observing System, which is located at every National Weather Service (NWS) field office to automatically collect wind, temperature, precipitation, and sky coverage data that are automatically assimilated into NOAA's numerical weather models; Advanced Weather Information Processing System, which is used at every NWS operational site to allow forecasters to collaboratively interact with other forecasters, data sources, and automated model output to generate weather warnings and graphical and text products; and NOAA's OceanReports Tool, which is an interactive geospatial toolkit that allows users to select from over 100 coastal and ocean databases to inform ocean and coastal planning.

Economy

As noted earlier, FHWA collects large amounts of geometric, traffic, and crash data in the Highway Safety Information System. FHWA has begun research on using machine learning to link and analyze these data. For example, the work will enable linking of crash reports with hospital records so that the severity of accidents can be better recorded and analyzed. Machine learning will allow DOT to ingest a large amount of safety-related data, leading to potential improvements in quality assurance and quality control of these data. The DOT Secure Data Commons is a cloud-based analytics platform that enables traffic engineers, researchers, and data scientists to access transportation-related datasets.[32] It provides a secure platform for sharing and collaborating on research, tools, algorithms, and analysis involving sensitive datasets using commercially available tools without needing to install tools or software locally. A current bottleneck is the slow process of manually labeling and cataloging written reports. FHWA is developing natural language processing tools for automatically labeling and cataloging Occupational Safety and Health Administration work zone accident and injury reports. For image data, FHWA is working with DOE national laboratory collaborators on a system that automatically removes personally identifiable information from images that contain people so that data can be released to researchers.

DOE data resources and AI tools are helping to improve domestic energy supplies. Current efforts harness multiple existing and streaming data sources to increase energy-economic efficiencies. For example, a signature library of component failure mechanisms is helping identify grid vulnerabilities to mitigate economic disruptions.

[32] https://its.dot.gov/data/secure/index.html

NTIA is building a wireless testbed for monitoring and characterizing spectrum usage, and it collaborated with other organizations to develop a standard data specification for spectrum data. In 2019, the system consists of six spectrum sensors distributed over a neighborhood-sized area that will continuously monitor a selected frequency band and store all collected data for use in research and analysis. NTIA is leading the development of a National Spectrum Strategy that, as directed by the 2018 Presidential Memorandum on Developing a Sustainable Spectrum Strategy for America's Future,[33] will build a secure, automated capability to facilitate assessments of spectrum use and expedite coordination of shared spectrum access among Federal and non-Federal spectrum stakeholders. Dynamic coordination of spectrum assignments is a complex task currently performed by human spectrum managers, and hence, is an opportunity for application of AI techniques. NTIA is exploring the use of AI for detection and classification of wireless signals, as well as for monitoring spectral activity for abnormal or suspicious behavior. Once operational, the spectrum monitoring database will be made publicly available to enable other researchers to apply ML and other AI techniques to these extremely large datasets.

NTIA also maintains publicly available collections of high-quality audio and video data for use by other researchers. Audio data were gathered to evaluate human perceptions of audio quality and have many samples in which the human voice is masked by background noise, such as sirens, saws, and gunshots. These datasets are available for researchers to apply AI techniques to extract speech from cluttered audio streams. NTIA also maintains the Consumer Digital Video Library (CDVL), a large collection of high-quality uncompressed video. CDVL contains video for which subjective tests have been performed with human subjects to rate the perceived quality of the video—a valuable resource for research use to train AI systems against a baseline of human perception.

Health

NIH goals include hosting or otherwise making available biomedical datasets, advancing computational infrastructure to support the biomedical research community, and supporting research programs that use AI technology for advancing biomedical knowledge. As a first step, NIH has developed the Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) program in partnership with commercial cloud service providers (see sidebar on previous page) to reduce technological and economic barriers to accessing and computing on large biomedical datasets to accelerate biomedical advances.

NIH	***Cloud Resources***
NIH's STRIDES Initiative is working with commercial cloud providers to reduce economic and technological barriers to accessing, and computing on, large biomedical datasets. A central tenet of the STRIDES Initiative is that data made available through these partnerships will incorporate standards endorsed by the biomedical research community to make data findable, accessible, interoperable, and reusable (FAIR). In its pilot phase, cloud service providers and NIH investigators are setting up cloud storage and services around test case datasets to develop principles, policies, and processes. The first of these datasets is the Sequence Read Archive (SRA),[34] which offers raw sequencing data and alignment information from high-throughput, next-generation sequencing platforms and technology. Through STRIDES, NIH is making available biological sequence data to the broader research community to enhance reproducibility and allow for new discoveries by comparing datasets. Additional datasets and services are expected to become available to the NIH-supported community over time, along with refined policies and procedures for the initiative.	

[33] https://www.whitehouse.gov/presidential-actions/presidential-memorandum-developing-sustainable-spectrum-strategy-americas-future/

[34] https://www.ncbi.nlm.nih.gov/sra

In 2018, NIH convened a two-day NIH Workshop on Artificial Intelligence in Medical Imaging to identify gaps and hurdles and to outline research roadmaps for advancing foundational and translational research to accelerate widespread clinical adoption and improved patient outcomes. Several roadblocks to disseminating machine intelligence applications in medical imaging were identified. A significant pressing issue was recognized as the absence of large (annotated) medical image datasets. Individual medical imaging datasets are relatively small (numbering in the hundreds to thousands) and reflect only local populations by lacking diversity/heterogeneity, rendering it difficult to develop universally useful and transferable algorithms that accurately recognize clinically relevant pathology. NIH is currently working with the national medical imaging community to address this problem through a program mirroring the NIH STRIDES initiative, described above.

A project led by the Office of the National Coordinator for Health Information Technology (ONC) in collaboration with the NIH National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK), funded by the Patient-Centered Outcomes Research Trust Fund, will capture lessons learned and best practices in the course of developing training datasets. This project will use the NIH STRIDES and Department of Veterans Affairs (VA) Data Commons cloud service models to leverage data from the U.S. Renal Data System which is co-sponsored by the NIH NIDDK, and the Centers for Medicare and Medicaid Services (CMS). Curated versions of CMS-collected data will be combined with electronic health record (EHR) data and insurance claims-based Veterans Health Administration (VHA) data to elicit novel insights into chronic and end-stage renal disease and which treatments are most effective according to patient characteristics and preferences. The project is part of ONC's interest in advancing standardized sharing of electronic health data via application programming interfaces, which will be critical to making large, high-quality datasets available for wider use and adoption of AI and machine learning in healthcare.

The Department of Veterans Affairs Data Commons (see sidebar) is being designed to bring VA research data and computing resources together with tools for enabling AI. Data includes the world's largest genomic repository via the Million Veteran Program linked to electronic health record-derived information as well as the VA Corporate Data Warehouse. A set of principles for the VA Data Commons was developed using VA strategic planning and veterans' considerations as inputs for initial piloting.[35] These inputs were also

VA	*VA Data Commons*

To empower research into veteran health needs, the Department of Veterans Affairs is designing a national VA Data Commons to integrate clinical, genomic, imaging, and other data from the country's largest integrated healthcare system with a scalable compute infrastructure.

The VA has the largest integrated healthcare system in the United States, with 1,243 healthcare facilities. The VA Corporate Data Warehouse has health information for over 22 million veterans. In addition, the Million Veteran Program (MVP) has genomic information for over 750,000 people. The VA Data Commons is bringing such datasets together into a scalable infrastructure for AI to address veterans' needs.

The principles developed for the VA Data Commons include that they contain VA data; access capability outside the VA; are collaborative with managed access; are secure, scalable, sustainable, usable, and designed for longevity; and are reproducible and reusable.

Use cases being explored include, as examples, Genome-Wide Association Studies (GWAS) leveraging MVP-based genomic data, traditional health services research such as examining drug cardiotoxicity across certain veteran populations, and study capabilities for nonprogrammers for running prespecified analysis workflows.

[35] https://www.research.va.gov/data/the_national_va_data_commons_pilot.pdf

mapped to other efforts and FAIR principles to do a gap analysis and update the list of principles accordingly.

HHS developed HealthData.gov as a site dedicated to making data discoverable and making valuable government data available to the public in hopes of better health outcomes for all. The data are collected and supplied from HHS agencies, including the Centers for Medicare and Medicaid Services, Centers for Disease Control and Prevention, Food and Drug Administration, and Agency for Healthcare Research and Quality, as well as from state partners. Topics include environmental health, medical devices, Medicare and Medicaid, social services, community health, mental health, and substance abuse.

FDA has developed a prototype to connect the FDA's entire infrastructure of critical data assets.

Justice & Security

IARPA proactively releases data from virtually all of its research programs, including datasets on satellite imagery, biometrics, and geopolitical events. Sometimes creating these datasets also requires creating new data annotation techniques. As an example, the DIVA program is building video analysis tools that automatically detect when people or vehicles perform specific actions. This requires efficiently annotating hours of footage for rare activities, which is a novel data annotation task. In addition to releasing hours of annotated data, IARPA also has released the tools and methods for performing this annotation so that others can use and improve them.

From an energy-security perspective, DOE shares AI resources publicly, to the extent practicable, through DOE and open-source platforms.

Strategy 6: Measure and evaluate AI technologies through benchmarks and standards.

In order to produce safe, effective, and socially beneficial AI, developers, users, and policymakers will need benchmarks and standards for evaluating the capabilities and risks of AI technologies. Federal agencies are playing a leading role in developing these new benchmarks and standards.

As the U.S. Government's hub for technology measurement and standards work, NIST is a significant AI research center in its own right and is well positioned to lead Federal efforts in AI standards-setting and benchmarking (see sidebar). Longtime research programs at NIST and a number of other Federal agencies in search and summarization, language understanding, language translation, biometrics pattern recognition, video analytics, new materials discovery, and robotics are both directly improving AI technologies and generating experience and data for AI benchmarking and standards-setting efforts.

NIST	Standards for Manufacturing

NIST's development of reference data and AI techniques has helped to create advanced materials for manufacturing. In support of the Materials Genome Initiative, NIST is expanding the reference data in its Joint Automated Repository for Various Integrated Simulations. This repository is designed to automate materials discovery using classical force-field, density functional theory, machine learning calculations, and experiments. All models, scripts, and postprocessing tools to train the models are available online,[36] and in less than two years they have been consulted more than 30,000 times.

NIST's creation of measurements and standards are also advancing robotics techniques. For three years, the NIST Agile Robotics for Industrial Applications Competition has offered cash prizes for robotics techniques that achieve agility in complex factory environments; the competition is evaluated by NIST-developed metrics.[37]

In August 2019, NIST released the report *U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools*[38] in response to the Executive Order (EO 13859; see page 1) directing Federal agencies to ensure that the Nation maintains its leadership position in AI. The plan identifies nine areas for AI standards: concepts and terminology, data and knowledge, human interactions, metrics, networking, performance testing and reporting methodology, safety, risk management, and trustworthiness.

Standards Research

DOE researchers are actively utilizing machine learning to aggregate enormous operational datasets in order to benchmark performance and identify areas where current predictive models fall short. The goal is to improve predictive models and reduce uncertainty in performance.

As part of its broader AI research agenda, NIST has increased its investment since 2017 in AI-related research and benchmarking programs, which are necessary for standards development. This work has included verification and validation of AI systems. For example, NIST's new FARSAIT program is already supporting several research projects related to standards and benchmarks, including projects to assess the performance of generative adversarial networks (a class of common AI methods) and AI-based robot systems, to improve detection and correction of accidental bias in AI systems and to measure the vulnerability of AI image-recognition tools to adversary attacks. NIST is coordinating these AI standards

[36] https://jarvis.nist.gov/

[37] https://challenge.gov/a/buzz/challenge/999/ideas/top

[38] https://www.nist.gov/sites/default/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

activities with other Federal agencies; for example, NIST is currently working with IARPA to test and validate tools developed for IARPA's TrojAI program.

Global Engagement with Standards Organizations

NIST also supports the development and adoption of international technical standards for AI (see sidebar below). NIST AI researchers participate in and provide leadership to key standards-setting organizations, including the International Organization for Standardization (ISO); the International Electrotechnical Commission (IEC); the American Society of Mechanical Engineers (ASME); the Institute of Electrical and Electronics Engineers (IEEE); the Data Mining Group; the Open Applications Group; and the American Society for Testing and Materials (ASTM) International. Within the ISO/IEC principal subcommittee of the Joint Technical Committee for AI standards, ISO/IEC JTC 1 SC 42, Artificial Intelligence, NIST serves as convener of the Big Data working group and is actively supporting other standards development activities.

DOC, DoD, NIST, NSF, & State Department	***International Standards Activities***
NIST staff have led the development of the IEEE Core Ontology for Robotics and Automation, in collaboration with over 175 members representing 23 countries from academia, industry, and government. The standard has been purchased over 2000 times and has spawned additional standards development for autonomous systems. NIST and DoD staff are also supporting Federal Government international engagements in the ISO/IEC's experts group focused on AI standards, JTC 1 SC 42, which NIST and DoD helped stand up through U.S. national body participation and support. JTC 1 SC 42 has working groups developing technical standards for foundational AI concepts, terminology, and framework; big data; trustworthiness; use cases and applications; and computational approaches and characteristics. Staff from NSF, DOC, and the U.S. Department of State were among a group of more than 50 international experts who developed the Organisation for Economic Cooperation and Development (OECD) Principles on Artificial Intelligence[39] that OECD adopted in May 2019. The experts group consisted of representatives of 20 governments as well as leaders from the business, labor, civil society, academic, and science communities.	

[39] https://www.oecd.org/going-digital/ai/principles/

Strategy 7: Better understand the national AI R&D workforce needs

Numerous Federal AI R&D programs aim to ensure that the American workforce is capable and ready to meet the rapidly increasing demand for AI professionals. These efforts align with the NTSC's 5-year plan for advancing science, technology, engineering, and mathematics (STEM) education.[40] Progress to date includes training AI researchers through several of the aforementioned AI R&D programs, as well as investigating the needs of the American AI R&D workforce and pursuing novel approaches to train such a workforce. Agencies support graduate students in AI-related research projects, introduce undergraduate students to big data and AI research opportunities, provide diverse pre-K–12 students with the fundamental skills and competencies associated with computational thinking to succeed in a digital world, and enable opportunities and support for entrepreneurs and communities to learn applications of AI technology. Several Federal programs have recently been created that directly support research on understanding the changing nature of work and predicting future workforce needs.

Training Future AI R&D Practitioners

Various Federal agencies are providing support for AI R&D training at all levels, from pre-K–12 to postgraduate and lifelong learning. NSF is supporting development of new instructional materials and teacher professional development enabling high-quality, rigorous, and engaging computer science education at the pre-K-12 levels (see sidebar).

NSF began prioritizing computational and data-enabled science and engineering in 2018 in a subset of awardees of its Graduate Research Fellowships Program. And NIJ has emphasized the application of innovative AI methods and tools to the vast criminal justice challenges as part of its Graduate Research Fellowship Program, which supports doctoral students engaged in applied research related to criminal justice.

Other Federal agencies are supporting prioritizing computational and data-enabled science and engineering in their graduate fellowship programs. For example, in 2018, DOE added a new track to its Computational Science Graduate Fellowship program. This track supports

NSF	***Computational Thinking Through Education***

Sustained NSF support over the last decade has laid the foundations for the development of new instructional materials and teacher professional development enabling high-quality, rigorous, and engaging computer science education at the pre-K–12 levels. For example, between 2009 and 2016, NSF funded the development of a framework for a new Advanced Placement® (AP) Computer Science Principles (CSP) exam, along with aligned curricula and associated teacher professional development. The first AP CSP exam was administered in May 2017; more than 50,000 students took the exam—the largest launch of any AP exam in the College Board's 60-year history.

Most importantly, the exam has resulted in a dramatic increase in the diversity of the students taking an AP computer science exam, with 183% more female students, 256% more Black/African American students, and 229% more Hispanic/Latinx students taking an AP computer science exam in 2018 as compared to 2016.[41] This effort illustrates key outcomes of Strategy 7, including the need for firm grounding in computational thinking through computer science education starting at the pre-K–12 levels, as well as the imperative to broaden the participation of groups traditionally underrepresented in computing and related fields.

[40] Committee on STEM Education of the National Science & Technology Council. *Charting a Course for Success: America's Strategy for STEM Education.* Dec. 2018; https://www.whitehouse.gov/wp-content/uploads/2018/12/STEM-Education-Strategic-Plan-2018.pdf

[41] https://reports.collegeboard.org/ap-program-results/ap-computer-science-growth

students who pursue advanced degrees in applied mathematics, computer science, or statistics, and it promotes effective use of high-performance systems, including in the areas of AI, ML, and deep learning.

Agencies also are supporting undergraduate and graduate students through standard AI research grants. For example, NIFA runs the Research and Extension Experiences for Undergraduates (REEU) program to promote experiential learning experiences for undergraduates. These projects provide students hands-on experience in research and agricultural extension projects to develop the skills necessary to join the agricultural workforce or pursue graduate studies in food and agricultural sciences. In 2018, REEU supported projects specifically in food and agriculture cyberinformatics and tools, and beginning in 2019, the program also promotes and supports projects that focus on applications of advanced AI methods to food and agriculture. Similarly, NIJ provides investigator-initiated technology R&D funding opportunities that are multidisciplinary in nature to foster innovative AI research solutions; these approaches enable training opportunities across various disciplines, including the behavioral, computer, physical, and social sciences.

The Changing Nature of Work

In 2018, NSF made its first awards for its Future of Work at the Human-Technology Frontier Big Idea. Supported projects employ the joint perspectives of computer science, engineering, education, and social science to (1) understand how technology is changing the nature of work, and (2) design technology that will help workers be prepared, productive, and engaged in the future workplace. The goal of the FW-HTF program is to advance understanding of how technology and people interact, distribute tasks, cooperate, and complement each other in different specific work contexts of societal importance, ranging from manual labor to knowledge work.

An important component of the AI education research portfolio is the development of technology to support lifelong learning and retraining for a future where some kinds of jobs disappear and other kinds are created at a rapid rate. An NSF program, Cyberlearning for Work at the Human-Technology Frontier, supports early-stage research specifically on innovative technology to improve STEM education (e.g., see sidebar, next page).

DOE is leveraging AI to enhance the productivity of engineers as they design advanced energy-related products and services, by supporting the development of ML-enhanced design tools that are intended to automatically generate candidate solutions to engineering problems given available inputs (e.g., sunlight, fuel); desired process outputs (e.g., electricity, thrust); and associated performance metrics (e.g., efficiency, cost). This design capability is expected to enable engineers to consider a wider range of novel solutions by helping them develop options outside of their personal experience bases and by freeing them from the need to construct and evaluate simulations of manually developed candidate solutions.

Other Federal agencies are pursuing approaches to enhance their workforces. For example, NOAA is using its Executive Level Research Council and Education Council to identify opportunities to build AI expertise across the agency and in the ocean, weather, marine transportation, fishing, energy, tourism and recreation, environment, and conservation sectors. In order to build AI expertise within the agency, NOAA sponsors employees to take courses on unmanned systems operations for bathymetric, ecosystem, and geographic surveys, and it encourages its employees to participate in workshops on AI and machine learning for ocean, fishery, weather, and conservation applications. NOAA's public outreach activities include sponsoring robotics workshops at aquaria and schools across the country to educate the public on AI applications for ocean exploration.

Similarly, the Census Bureau has created the Statistical Data Modernization (SDM) project to bring its workforce, operations, and technologies up to date and set the standard for statistical agencies in

today's data-driven society. The workforce transformation component of SDM is enabling the hiring of new data scientists with expertise in the application of new methods and analytics, including the use of AI methods and tools to process and analyze big data. The workforce transformation also is addressing the upskilling of the current data scientist workforce. The SDM team is working with the Office of Personnel Management and other agencies to establish a data scientist series that will encourage the hiring of talent in AI and related technologies across the Federal Government.

NSF	***Information Technology (IT) and the U.S. Workforce***

In 2014, NSF funded the National Academies of Sciences, Engineering, and Medicine (NASEM) to examine current and possible future impacts on the U.S. workforce of emerging information and communication technologies. The charge to NASEM was framed broadly: assess the many dimensions of the evolving relationship between technology and work and set out a research agenda. NASEM convened a 13-member committee that in 2016 ran an information-gathering workshop and issued a report in 2017, *Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?*, articulating the current state, trends, and possible futures of technology and work as follows:[42]

1. *Advances in IT are far from over, and some of the biggest improvements in areas like AI are likely still to come. Improvements are expected in some areas and entirely new capabilities may emerge in others.*
2. *These advances in technology will result in automation of some jobs, augmentation of workers' abilities to perform others, and the creation of still others. The ultimate effects of information technology are determined not just by technical capabilities, but also by how the technology is used and how individuals, organizations, and policymakers prepare for or respond to associated shifts in the economic or social landscape.*
3. *The recent increase in income inequality in the United States is due to multiple forces, but these include advances in IT and the patterns of its diffusion, globalization, and economic policy.*
4. *IT is enabling new work relationships, including a new form of on-demand employment. Although current digital platforms for on-demand work directly involve less than 1 percent of the workforce, they display significant growth potential.*
5. *As IT continues to complement or substitute for many work tasks, workers will require skills that increasingly emphasize creativity, adaptability, and interpersonal skills over routine information processing and manual tasks. The education system will need to adapt to prepare individuals for the changing labor market. At the same time, recent IT advances offer new and potentially more widely accessible ways to access education.*
6. *Policymakers and researchers would benefit significantly from better understanding of evolving IT options and their implications for the workforce. In particular, (1) sustained, integrated, multidisciplinary research and (2) improved, ongoing tracking of workforce and technology developments would be of great value for informing public policies, organizational choices, and education and training strategies.*

These findings, along with recommendations for a research and education agenda focused on work at the human-technology frontier, have influenced Federal agency investments in recent years, including NSF's FW-HTF Big Idea and related programs enabling lifelong learning.[43]

[42] https://www.nap.edu/catalog/24649/information-technology-and-the-us-workforce-where-are-we-and, pp. 9-10.
[43] https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504984

Strategy 8: Expand public-private partnerships to accelerate advances in AI

While private-sector innovation in AI R&D is proceeding at a significant pace, the vast majority of industry R&D is focused on near-term applications. The Federal Government, on the other hand, invests in high-risk, long-term research and is uniquely positioned to provide leadership and facilitate collaboration. Updated mechanisms for public-private partnerships are needed to leverage private sector energy to address long-term research and public sector challenges. Multiple Federal agencies are pressing forward to find new ways for the public and private sectors to collaborate that are appropriate to the fast-moving realm of AI R&D. Their efforts fall into two broad categories: easing access to data and computing resources, and jointly funding research. These efforts in public-private partnerships are a valuable start, but challenges remain. One important step would be for agencies that have made progress to share their best practices for building public-private partnerships for AI R&D.

Science & Engineering

NSF engages in public-private partnerships to advance research in numerous areas, including AI. In 2018, NSF encouraged proposals for Early-Concept Grants for Exploratory Research (EAGERs) on understanding and mitigating social challenges arising from AI technology, working with the Partnership on AI, a nonprofit organization with public- and private-sector members. In 2019, NSF and a major U.S. technology company committed to jointly fund research on fairness in AI. Although industry provides part of the funding, award decisions are made independently by NSF through its standard merit review processes. NSF has also partnered with another leading U.S. technology company to jointly support multiple funding opportunities over the course of five years that have led to a number of advances in AI (see sidebar).

NSF has responded to the compute-intensive demands of artificial intelligence and machine learning through use of public-private partnerships. In 2017 and 2018, NSF partnered with four major cloud computing vendors to make cloud computing

NSF	***AI-Based Reconstruction Algorithms for Imaging***

Jointly funded by NSF and an industry partner through the Partnership on Visual and Experiential Computing, a team of researchers has made major advances in light field imaging by combining algorithmic and software innovations, with developments in physical optics system design.[44] In particular, light field-enabled monocentric optics and AI-based reconstruction algorithms are making possible spatially varying digital foci for complex and wide-field-of-view scenes and 3D imaging capabilities. Moreover, this research has led to the emerging field of *deep optics.* Whereas conventional digital cameras record raw sensor data and algorithms then process it to compute an image, deep optics allows for the co-design of optimized optical systems and AI algorithms. This approach unlocks the full potential of optimized application-domain-specific camera systems for light field, multispectral, high dynamic range, and other imaging modalities. Using this AI-driven approach to imaging, the group demonstrated significant advances to the perception capabilities in robotic vision, remote sensing, medical imaging, autonomous vehicles, and defense. In parallel, innovations in the optics hardware design led to the formation of an imaging startup company that is exploring innovations at the crossroads of AI, virtual reality (VR) and visual computing, for example, in drones and 360-degree VR applications.

[44] VEC: Small: Collaborative Research: Wide Field of View Monocentric Computational Light Field Imaging, https://www.nsf.gov/awardsearch/showAward?AWD_ID=1539131.

resources available to academic researchers through its BIGDATA program. Building upon the success of this effort and a 2018 conference, Enabling Computer and Information Science and Engineering Research and Education in the Cloud,[45] NSF/CISE announced a new solicitation, Cloud Access,[46] to further explore the role of the commercial cloud in CISE-funded research. Cloud Access aims to support a liaison entity linking the CISE community with public cloud computing providers to meet the cloud computing needs of the CISE research and education community. This solicitation, awarded in September 2019 to a trio of universities, will provide access to cloud resources and other services and strategic technical guidance and training in using the cloud. After evaluating the success of such an entity, NSF will consider expanding the program to meet the needs of the science and engineering research community more broadly.

Public-private partnerships are at the core of the Department of Energy mission. Through the High Performance Computing for Energy Innovation Program (HPC4EI),[47] DOE pairs National Laboratory scientists with industry partners to apply HPC technology and improve energy utilization in the manufacturing sector. HPC for Manufacturing aims to advance innovative clean energy technologies, reduce energy and resource consumption, and infuse advanced computing expertise and technology into the manufacturing industry. HPC for Materials focuses on material performance predictions at multiple phases/scales, data interpretation and real-time analytics, and integration of multi-physics of HPC codes to simulate material behavior. DOE is using public-private partnerships, including those established through the HPC4EI program, as a means of exploring potentially disruptive AI technology to meet increasing mission needs. For example, its recently formed partnership with a Silicon Valley startup will investigate a new class of computer systems that has the potential to dramatically accelerate AI workloads.[48]

The NIST National Cybersecurity Center of Excellence[49] is the Department of Commerce's only federally funded R&D center. The center collaborates with academia and industry to provide real-world, standards-based cybersecurity capabilities that address business needs. The NIST NCCoE is currently working with several industry partners on projects documenting terminologies applicable to securing AI systems and developing a taxonomy that addresses attacks and defenses to AI systems and technologies. As is common with NCCoE projects, the end result will be publicly available guidance. Future planned NIST NCCoE projects related to AI in FY2020 include an effort to collaborate with industry partners in a laboratory environment to demonstrate trustworthy AI capabilities.

NOAA is engaged in a variety of public-private partnerships through cooperative research and development agreements, contracts, and other agreements to advance AI applications to meet all aspects of NOAA's mission. Examples include working with commercial cloud computing providers to develop AI and machine learning capabilities to assess marine mammal and fish stock abundance from aerial and ship-installed image collection systems and underwater acoustic recordings; employing robotics for deep sea exploration; providing automated control of NOAA's 18 environmental satellites with commercial providers; using machine learning to analyze satellite imagery for severe weather detection and prediction, oil spill and hazardous material trajectories, wildfire detection and

[45] https://dl.acm.org/citation.cfm?id=3233928.

[46] Enabling Access to Cloud Computing Resources for CISE Research and Education (Cloud Access); https://www.nsf.gov/pubs/2019/nsf19510/nsf19510.htm.

[47] https://hpc4energyinnovation.llnl.gov

[48] https://www.businesswire.com/news/home/20190917005356/en/Department-Energy-Cerebras-Systems-Partner-Accelerate-Science .

[49] https://www.nccoe.nist.gov/

movement, ecosystem health, and detection of illegal fishing activity; and leveraging private-sector data such as crowd-sourced bathymetry for data-mining ocean and coastal big datasets to allow users to efficiently determine siting feasibility for maritime applications such as aquaculture, renewable and nonrenewable energy development, critical minerals extraction, fishing, coastal development, tourism and recreation, national and homeland security, and environmental conservation.

Defense

DoD is accelerating department-wide industry collaboration in AI. It established the Joint Artificial Intelligence Center (JAIC) in June 2018 to accelerate the delivery of AI-enabled capabilities focusing on a common set of AI standards, tools, shared data, reusable technology, processes, and expertise for all of DoD, including through evolving partnerships with stakeholders such as industry, other Federal agencies, and technology vendors. JAIC is guiding execution of National Mission Initiatives.

Alongside JAIC, one of the first deployed initiatives of AI comes from Project Maven (see sidebar), which focuses on computer vision aboard unmanned aerial vehicles, often used for humanitarian assistance and in disaster relief situations. This program is a key example of public-private partnerships in technology. Over 130 small and large companies participated in a 2017 Project Maven industry day, demonstrating interest in participating in operational AI development with the DoD. DARPA also works closely with industry through numerous public-private partnerships to jointly fund research in AI for long-term breakthroughs.

DoD	Project Maven Partners with Industry

The Department of Defense has focused its algorithmic efforts on iteration and fielding with commercial partnerships. Starting in late 2017, the Undersecretary of Defense for Intelligence's Algorithmic Warfare Cross-Functional Team, also known as Project Maven, took a commercial approach to fast-moving algorithm research and development. Within six months, the AI R&D team was able to procure, refine, and deploy artificially intelligent algorithms to warzones to complement and work with human processes.

The imagery classification algorithms that Project Maven works with are almost entirely licensed from top-tier American technology companies and startups. The team takes a commercial approach because of the wealth of technical talent in private industry, matched with a rigorous security review. Project Maven focuses on testing and evaluation of external technologies to keep abreast of the fast pace of global advances in machine learning rather than on internal development. This method has offered an opportunity to build trust and mutual standards between Silicon Valley companies and the Pentagon.

In 2019, two years into its existence, Project Maven's computer vision R&D has introduced technological advances that reduce civilian casualties in warfare. Furthermore, the team has expanded to numerous lines of effort, training commercial algorithms at various classification levels using cloud infrastructure. One of the uses of computer vision algorithms is for humanitarian assistance and disaster response, to reduce the time associated with search and discovery, resource allocation decisions, and rescue operations to save lives.

Economy

NIFA's competitive grants program Food and Agriculture Cyberinformatics and Tools encourages university-based research as well as public and private partnerships. FACT encourages robust representation of public and commercial entities from relevant disciplines that bring together research, education, and/or extension communities working on multiple aspects of a data-supported problem into one multidisciplinary network to collectively advance the field. In addition, FACT supports efforts that will develop and maintain effective networking strategies for connecting the public and private sectors, students, and the media to foster interest, economic development, and societal engagement.

Health

As part of its innovation efforts around digital health, HHS's Office of the Chief Technology Officer, together with the Presidential Innovation Fellows, piloted a new "TOP (The Opportunity Project) Health Tech Sprint" initiative with new models of data release for AI training and testing and a voluntary incentivization framework for an AI ecosystem between public and private entities.[51] Eleven collaborative teams including technologists, issue experts, and community leaders investigated use of digital tools built with Federal data and AI technologies to improve clinical trials, experimental therapies, and data-driven solutions for complex challenges ranging from cancer to Lyme and tick-borne disease.

The VA launched its own AI Tech Sprint (see sidebar) that has engaged over a dozen companies to pilot a novel approach to private-public partnership. In addition, the VA AI Tech Sprint uses the VA-HHS-GSA AI-Able Data Ecosystem (chosen as one of the top public-sector innovations of 2019 by Government Innovation Awards) for qualitatively measuring and incentivizing discoverability and usability of AI models and data from both Federal and non-Federal perspectives. In combining these, the AI-Able Data Ecosystem promotes a nimble collaborative framework by building a data linkage between data producers and AI/model creators.

VA	*AI Tech Sprint*

Modeled broadly after the HHS TOP Health Tech Sprint, the private–public VA Tech Sprint partnership with a U.S.–international health technology company jointly published an article in *Nature* in August 2019 on a deep learning approach for prediction of acute kidney injury.[50] The work is an example of collaborative agreement on large datasets, ones that can take years to develop.

The AI Tech Sprint has been piloting a novel, lightweight mechanism for establishing early-stage collaborations that can be potentially transformed into partnerships. A process has been implemented whereby small real and synthetic datasets in the VA format for health information are developed using simplified Data Use Agreements that can be executed in days. This data can then be provided directly to non-Federal entities to enable these collaborators to create tools before, or in parallel with, drawing up of more complex contracts for larger, richer datasets. Thus, once the additional contracts are executed, the tools have already been developed for the specific VA data types and AI applications, greatly speeding potential development and deployment to help veterans. This is being done as part of a Parallel Pathway to Potential Partnerships, whereby parties interested in donating funds or R&D, and other potential collaborations, can apply prior to more extensive contractual commitments.

Beyond linking companies to Federal data for creating novel tools, the AI Tech Sprint has fostered new collaborations among non-Federal partners around an ecosystem of Federal data. In one such example, applications have been developed that can empower millions of veterans and others to find clinical trials and experimental therapeutics directly from the VA and the Centers for Medicare and Medicaid Services, based on their own health information.

Justice & Security

The DHS Science and Technology Directorate brings workable solutions, including AI and machine learning, to data analytics challenges facing the homeland security enterprise. Meeting these needs requires overseeing the process of transferring technology from the laboratory to the marketplace, which in turn necessitates partnerships with industry and academia. DHS's Data Analytics Technology Center regularly establishes cooperative research and development agreements with a range of industry labs to understand the state of commercial practice. DHS's Silicon Valley Innovation Program

[50] https://www.ncbi.nlm.nih.gov/pubmed/31367026

[51] https://www.hhs.gov/cto/blog/2019/1/17/top-health-tech-sprint-unleashes-the-power-of-open-data-and-ai.html

expands the Directorate's reach to find new technologies that strengthen national security with the goal of reshaping how government, entrepreneurs, and industry work together.[52] Based in California's Silicon Valley, the program aims to harness the commercial R&D ecosystem for technologies with government applications and to co-invest in and accelerate technology transition to market.

In support of its criminal justice and national security missions, the FBI is rapidly deploying advanced analytic methods to support operations. With the growth in data collected and available and the advances in low-cost computer power and storage, the FBI collaborates with industry and academia to ensure that research is conducted to meet its unique needs, for example, by applying machine learning and augmented intelligence to support FBI investigators and analysts. These efforts include working with venture capital to partner with other government agencies and industry in order to develop solutions in the AI space.

[52] https://www.dhs.gov/science-and-technology/svip

Conclusion and Next Steps

The Nation's AI R&D enterprise is strong. As defined in the President's Executive Order on AI, [53] U.S. R&D and innovation in AI is a critical component of an overall national strategy to maintain American global leadership in AI. It is incumbent upon the United States to continue to leverage its technical strengths and support early-stage research to drive technological breakthroughs in AI.

This 2016–2019 AI Progress Report describes some of the significant R&D progress in AI made over the last three years as a result of Federal investments. The 2016–2019 Progress Report includes specific examples of recent Federal agency programs that are aligned with the eight strategies described in the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* published in June 2019. While not exhaustive, the Federal programs and activities described constitute a comprehensive set of investments and accomplishments that are working to further the Nation's leadership in AI R&D.

This 2016–2019 Progress Report is the first periodic compilation of the progress of Federal programs and activities in meeting the goals outlined in *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update.* This progress is just part of the overall progress being made in the context of the Nation's AI strategy—the American AI Initiative. Collectively, the investments described in this report reflect how the Federal Government is leveraging and improving America's AI capabilities through R&D to ensure that those capabilities produce increased prosperity, safety, security, and quality of life for the American people for decades to come. Recognizing that Federal programs and activities must continually evolve to remain at the frontiers of AI R&D, future versions of this Progress Report on Federal R&D for Advanced Artificial Intelligence Technologies will be compiled to document ongoing progress in Federal AI R&D.

[53] https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/

THE WHITE HOUSE