

University of Pennsylvania Input to National Security Commission on Artificial Intelligence

WHITE PAPER

FEBRUARY 2020

University of Pennsylvania (Penn) Working Group Participants

Dawn A. Bonnell, Vice Provost for Research

Michael J. Borda, Director of Research Initiatives, Office of the Vice Provost for Research

Kostas Daniilidis, Ruth Yalom Stone Professor of Computer and Information Science

Ezekiel J. Emanuel, Vice Provost for Global Initiatives; Chair, Department of Medical Ethics and Health Policy

Amy E. Gadsden, Executive Director, Penn Global; Executive Director, Penn China Initiatives

John Gans, Director of Communications and Research, Perry World House

Lee R. Goldberg, Professor of Medicine

Michael C. Horowitz, Interim Director, Perry World House; Professor of Political Science

Lauren Kahn, Research Fellow, Perry World House

Vijay Kumar, Nemirovsky Family Dean of Penn Engineering; Professor of Mechanical Engineering and Applied Mechanics, Computer and Information Science, and Electrical and Systems Engineering

Jason H. Moore, Edward Rose Professor of Informatics; Director, Penn Institute for Biomedical Informatics

Scott Moore, Director, Penn Global China Program

Michael S. Parmacek, Frank Wister Thomas Professor of Medicine

Dan Roth, Eduardo D. Glandt Distinguished Professor of Computer and Information Science

Christian Ruhl, Program Associate for Global Order, Perry World House

Mitchell D. Schnall, Eugene P. Pendergrass Professor of Radiology; Chair, Department of Radiology

Jianbo Shi, Professor of Computer and Information Science

Special thanks to Lauren Kahn and Christian Ruhl for serving as the principal drafters of this report.

This report is the sole responsibility of the authors. It does not represent the positions, policies or opinions of Penn Global, Perry World House, the University of Pennsylvania, or the National Security Commission on Artificial Intelligence.

Executive Summary

In 2019, the U.S. government organized itself more clearly to ensure continued American leadership in artificial intelligence (AI). The Trump administration's AI strategy, outlined in a February 2019 [Executive Order](#), articulates the importance of AI for economic and national security leadership. More is necessary, however, to [sustain American leadership](#) moving forward, including new resources, incentives for partnerships, and new organizational vehicles to encourage AI breakthroughs.

In summer 2019, following discussions with the National Security Commission on Artificial Intelligence (NSCAI), and under the vice provost for global initiatives and Perry World House, the University of Pennsylvania (Penn) convened experts from its schools of engineering, law, and medicine; the political science department; and other disciplines to discuss potential recommendations for the commission. Given the [mandate](#) of the NSCAI, the Penn experts focused on two primary questions:

- 1) How should the United States use education, training, and workforce development to strengthen its AI talent base?
- 2) What institutional or regulatory reforms does the United States need to maximize its competitiveness in AI research and development?

Broadly, the Penn analysis posits that continued American leadership in AI in the 21st century will require a focus on developing people. Superior projects and partnerships require people with specific AI knowledge and skills. More and better AI talent will not just mean realizing the next breakthrough but also deploying technology to power improvement at multiple levels and across industries. In addition, the United States needs to attract people dedicated to the nation's interests, not just the bottom line.

Although neither national security nor technological research are zero-sum, the nature of preserving national leadership requires the U.S. government to be clear-eyed about structuring these programs. Up to this point, the United States, as well as institutions like Penn, have mostly focused on pursuing and promoting open research models that support not only America's interests but all of humanity's. That does not need to change, but to avoid the unnecessary leaking of research funding and talent, Penn offers a number of recommendations to steer investments and programming in ways that will benefit the United States.

In the following pages, the authors present two goals and eleven guiding ideas to address the primary questions above.

Goal 1: People First—Strengthen America’s AI Talent Base and Overall Research Leadership

- a) Initiate AI training and literacy at the secondary-school level.
- b) Rapidly expand master’s and doctoral programs using online and other nontraditional—but easily scalable—educational and training methods.
- c) Develop special recruiting programs and international collaborations to maximize access to the global AI talent base.
- d) Develop translational AI research capability, including interdisciplinary AI fellowships.
- e) Create tech ecosystems, not one-off capabilities.
- f) Incentivize broad AI research and adoption.

Goal 2: Ensure AI Leadership by Building Partnerships for an AI-Enabled Future

- g) Develop public-private research partnerships to bridge the academia-industry divide.
- h) Replicate successful interagency research models.
- i) Make long-term investments in people, not just projects.
- j) Ensure that AI expertise exists at senior levels of government.
- k) Work with universities and civil society to foster a national dialogue on AI ethics.

Goal 1: People First—Strengthen America’s AI Talent Base and Overall Research Leadership

a) *Initiate AI training and literacy at the secondary-school level.*

AI workforce training should start **well before undergraduate education through programs at the secondary-school level**. First, such training will begin growing a pipeline of domestic talent for AI development, ensuring a long-term supply of software engineers, data scientists, and workers who can develop and use innovative AI applications. Second, such training will build AI literacy in the large majority of the national workforce that does not pursue tertiary education. In 2017, only 35.7 percent of Americans had a bachelor’s degree or more; high dropout rates unresponsive to tuition price changes [suggest that it may be difficult to increase this number](#).

Secondary-level AI education would ensure that all parts of the American workforce could contribute to the AI revolution. The biggest leaps the United States has made in terms of relative power—economic, military, and technological—occurred when government engaged the American people toward a mission. Giving those in the United States the opportunity and tools to contribute will help do just that.

Several innovative pilot projects demonstrate pathways to increase AI training in high school. Early programs have partnered high-schoolers with universities, such as [Carnegie Mellon University](#) and [Stanford’s residential summer program in AI](#). The umbrella organization AI4ALL, which hosts several such programs, has a [track record](#) of providing AI training to high-school-level American citizens and lawful permanent residents, particularly from disadvantaged communities. [Research](#) suggests that investments in children born into such circumstances may have high returns for innovation; high childhood exposure to inventors may increase the likelihood of a child becoming an inventor by up to ten times.

Competing with [Chinese investments in secondary-level AI education](#) may require broader programs. For instance, investments in computer infrastructure and course materials in public schools would provide the basic foundations for children to learn about AI. Moreover, most existing U.S. AI-training programs are designed as college preparatory training for high-school juniors and seniors. Full preparation for an AI-based national economy will require not only people to develop new algorithms, but also people to use, apply, and interact with narrow AI in the workplace. Such training could include investments in [career technical education \(CTE\), which has proved to have high returns on investment](#).

Recommendation: Give the Department of Education’s Office of Innovation and Improvement a significant budget increase for secondary-level AI initiatives.

Penn Input to National Security Commission on AI

Recommendation: Increase the current \$1.3 billion of Department of Education [CTE State Grants](#) to \$2.6 billion to jump-start AI training in the civilian workforce.

Recommendation: Double the Department of Education's Office of Innovation and Improvement's budget for "Creating 21st Century Learners and Coders" from [\\$1.2 billion](#) to \$2.4 billion.

Penn Input to National Security Commission on AI

b) *Rapidly expand master's and doctoral programs using online and other nontraditional but scalable methods.*

Colleges and universities confer only 27,000 master's degrees in the broad area of computer science in the United States each year. While this number is doubling on a five-year time scale, this growth is nowhere near what is necessary for American leadership in computer science more generally, and very few of these degrees specialize in AI.

The University of Pennsylvania is planning to invest in on-campus master's programs, on-campus certificate programs, and other online offerings. In January 2019, Penn launched the very first Ivy League online master's program in computer science. This master's program allows Penn faculty to teach courses traditionally limited by classroom size to more than 600 students, resulting in a program size of 1,500 students. Penn also launched an on-campus certificate program in cybersecurity in fall 2019. The university projects that it will reach over 10,000 students in the next three years.

Expanding offerings in AI over the next five years at the University of Pennsylvania alone will **impact 50,000 students**. The goal is to make Penn a leader in AI training, growing its reach by 100 times. Central to the Penn plan is offering modular, customized programs that will be of interest to learners who want to be generalists or application domain experts and to learners who want credentials or to develop a specialized skill set. Students can combine the modules to learn foundational AI techniques, specific techniques for application domains (e.g., the AI a scientist needs to know for genetics research or for autonomy in drones, unmanned aerial vehicles, etc.), or general-purpose tools (e.g., in-depth with Apache Spark and its Machine Learning Library).

At the postsecondary level, investments in nontraditional methods to expand master's and doctoral programs in AI would help to rapidly develop American talent: 42 percent of data-science jobs require graduate degrees, and demand for AI and machine learning is far outstripping supply. Moreover, master's and doctoral programs may be easier to expand than undergraduate programs—as the former require fewer resources—and, because they focus on skills development, can more readily be delivered through online courses. Relatively small investments in online postgraduate education could therefore have large impacts on the development of talent.

In addition to Penn's experimentation, another possibility for rapid growth in training the future AI workforce comes from Georgia Tech's online master's degree in computer science, which has expanded the number as well as the type of students without decreasing in-person enrollment. The online cohort of this model also performed higher on tests than the in-person students. The online program attracted a significantly higher share of American citizens: 70 percent, compared with 8 percent for the in-person version of the course. This finding should help ease concerns about international free-riding on U.S. educational investments.

Penn Input to National Security Commission on AI

Funded with an initial \$4 million in outside start-up funds by AT&T, this program alone is estimated to increase U.S. production of master's degrees in computer science by [7 percent per year](#). At this rate, a \$607 million investment could double the number of master's degrees in computer science in the United States per year.

Recommendation: Investigate and develop best practices for rapid expansion of online AI training based on the Penn and Georgia Tech models, along with other universities.

Penn Input to National Security Commission on AI

c) *Develop special recruiting programs and international collaborations to maximize access to the global AI talent base.*

The investments in the development of domestic AI talent will likely take years or decades to bear fruit. For the United States—or any other country—to remain ahead in these critical technologies, it is essential to attract talent from around the world. Oren Etzioni, the head of the Allen Institute for Artificial Intelligence, [has noted](#) that two-thirds of the Institute’s research scientists are immigrants. In 2016, [almost 60 percent of all AI degrees](#) awarded in the United States were to nonresident aliens. This kind of approach mirrors one used by [countries such as Canada](#), whose 2017 Pan-Canadian Artificial Intelligence Strategy emphasizes attracting talent and has to date successfully recruited well-known AI researchers to Canadian universities and research institutes.

Leadership in AI will not just require a stronger domestic AI workforce but international cooperation as well. Increased international scientific cooperation would allow the United States to capitalize on the expertise of those researchers who are unwilling or unable to permanently relocate to the United States. A large-scale [mission-based](#) project for AI, modeled on the Human Genome Project (HGP), for instance, may result in useful breakthroughs (as well as international goodwill). The nature of AI research allows for broader resource sharing and therefore stronger collaboration than was possible with the HGP: Computing power and data access, unlike HGP wet lab equipment, is easier to share and access from anywhere in the world. Such a project may help to secure the U.S. lead in AI just like the HGP-seeded large biotechnology industries in the United States. Furthermore, unlike naturally occurring gene sequences, [which the U.S. Supreme Court has ruled unpatentable](#), the ownership of AI outputs may be easier to patent.

The United States should make it easier for high-skilled immigrants in AI-related fields to obtain work visas, including [J-1, H-1B, and L-1 visas](#). Penn bases this recommendation on the well-established link between immigration and innovation. A one-percentage-point increase in the immigration of advanced-degree holders is estimated to lead to a [27 percent increase in patenting per capita](#). On the flip side, studies of previous immigration quotas suggest that a decrease of 10 percent in new immigrants [decreases patent applications](#) by 0.5 percent per year.

Worries that H-1B visa holders take jobs from native-born or naturalized Americans are misplaced. There are more jobs than talented applicants in the United States. Moreover, increases in the H-1B visa cap correlate with an increase in patented inventions [without crowding out native scientists and engineers](#). A targeted [fast-track visa option](#) could prioritize immigrants with AI skills and data science in the H-1B visa process.

For context, the [fiscal year \(FY\) 2020 H-1B visa cap](#) was 65,000, with a 20,000 advanced-degree exemption.

Penn Input to National Security Commission on AI

A path to citizenship for those with skills in AI, including those with master's and doctoral degrees in relevant disciplines, would allow the United States to retain this talent and seed long-term growth in AI innovation. Lack of action on this front could lead to a brain drain to countries like China and Canada—[62 percent](#) of polled Canadian tech companies said they saw a spike in applications from the United States in 2017.

Recommendation: Create a start-up visa program or a [specialist AI visa](#) program. An improved [immigrant entrepreneur visa](#) would allow those with innovative ideas in AI, but without employee sponsorship, to start AI companies in the United States and pursue permanent residency to ensure their innovations remain in the United States.

Recommendation: The United States needs to look internationally and develop special recruiting programs. These may include fast-track visa options for skilled immigrants in AI, a redesigned “Start-Up Visa” program, a possible “AI Visa,” simpler paths to citizenship for highly skilled AI immigrants, and U.S.-funded large-scale international collaborations.

Recommendation: Leverage foreign talent through government incentives for U.S. institutions to build stronger partnerships with foreign AI researchers and research institutes in allied and like-minded countries. These partnerships could involve joint or matching funding on behalf of the United States and host country governments, or fellowships that would facilitate joint research projects and exchanges. This kind of approach mirrors a proposal in the U.S. Department of Defense’s AI Strategy to expand personnel exchanges with allied militaries in the AI field.

Penn Input to National Security Commission on AI

d) Create tech ecosystems, not one-off capabilities.

In addition to AI-literate and -capable talent, America's leadership and competitiveness in AI should be thought of in terms of creating multiple tech ecosystems. AI is a collection of algorithmic tools that cannot execute in the absence of **advanced computational and device infrastructure**, spanning supercomputer resources to the data-generating devices of the Internet of Things. The investments required to maintain leadership in AI will fail if they are competing with investments required to maintain global leadership in supercomputing, to support the Quantum Information Systems Initiative, and to advance device technology and integration. These are synergistic. All must flourish to ensure robust AI ecosystems.

Moreover, there is a fundamental gap between computing systems that perform brain-like functions and the brain itself. Current neuromorphic brain-like architectures have not achieved brain-like functionality and have only succeeded in providing point solutions to very well-specified problems (e.g., object recognition, speech-to-text translation, playing games, or taking tests).

The United States needs to invent new **bio-inspired architectures** (beyond graphics-processing units, tensor-processing units, and complementary metal-oxide-semiconductor accelerators) with collocated memory and processing and new software architectures so that scientists can develop truly efficient neural fabrics. Indeed, an important aspect of this work is understanding the complexity of the human brain and why combining existing U.S. strengths in neuroscience, bioengineering, and network science gives the United States a unique advantage.

Recommendation: Fund basic research and education on the human-machine interface, including bioengineering. These interfaces could enable new algorithms and frameworks to create intelligent systems that will interact with humans and augment human performance.

Penn Input to National Security Commission on AI

e) *Develop translational AI research capability, including interdisciplinary AI fellowships.*

Translational research is essential in the AI space because AI is a tool, not an end in itself. Effective applications of AI are likely to occur by linking basic research with applications in areas such as health care, robotics, business, national security, or elsewhere. Therefore, efforts to build AI capacity must consider how AI will connect to real-world problems, and their solutions and funding should target opportunities that will promote AI knowledge and further America's advantage in critical research areas.

For example, the United States is well-positioned to develop top translational research capability in **biomedical applications of AI**. AI is expected to [revolutionize all aspects of health care](#). Moreover, health-care spending was [17.9 percent of U.S. GDP](#) in 2017—the [highest percentage of any country in the Organisation for Economic Development and Co-operation](#)—meaning that AI's health-care applications will have profound effects on the national economy and on America's standing in the world. Taking advantage of the large budget and existing institutional strength of the National Institutes of Health (NIH), new AI fellowships could be housed in a revived [interdisciplinary research](#) initiative of the NIH Common Fund. This initiative lost most of its \$45 million funding in [FY 2012](#); restoring this amount for translational research in AI would require an increase of 8.4 [percent](#) to the [FY 2020 budget](#), to \$578 million. An interdisciplinary approach to translational research capability in the life sciences could also benefit fundamental research in machine learning, as deep learning, neural networks, brain emulation, and many other AI methods link closely to neuroscience and cognitive science.

Recommendation: Develop translational research capability with a \$45 million increase to the NIH Common Fund's interdisciplinary research budget, designated for new grantmaking funds and new interdisciplinary AI fellowships. This will allow the United States to maintain the lead in both the development and the creative application of AI in the most promising areas.

Penn Input to National Security Commission on AI

f) *Incentivize broad AI research and adoption.*

Beyond health care and the biomedical field, effective applications of AI will necessarily require **connecting AI knowledge with area-specific knowledge in fields such as mechanical engineering, business, law, and politics**. Given that [46 percent](#) of federal funding for civilian science already goes to the NIH, it is critical that the NIH funding increases recommended above not crowd out other AI funding. Funding beyond biomedical applications of AI will have significant spillover effects in private-sector innovation; [one study of federal grantmaking activity](#) suggests that every \$10 million of funding could lead to an additional 3.1 patents, mostly in translational applications that differ from those envisioned in the original grants.

Several programs within the National Science Foundation (NSF) could house innovative fellowships. The NSF overall received only [\\$5.6 billion](#) in federal funding for research in 2017. The \$30 million committed to the NSF “big idea,” [Harnessing the Data Revolution, in FY 2019](#) is an excellent start, but more is necessary.

Recommendation: Commit \$12 billion a year to double non-biomedical research funding in strategic AI-related areas.

Recommendation: Increase NSF funding for AI in several areas:

- **Expand funding for NSF [Graduate Research Fellowships](#) for AI and data-science programs to encourage interdisciplinary research and to prioritize those projects that apply AI to other areas, such as robotics and material science.**
- **An NSF interdisciplinary research initiative should prioritize AI as an “[area of national importance](#)” in its grantmaking.**
- **Expand the NSF Research Traineeship to fund new avenues for graduate student training in “[high-priority interdisciplinary or convergent research areas](#).”**

Recommendation: Explore other AI-focused interdisciplinary fellowships throughout the U.S. government. For example, a fellowship under the Justice Department’s [Justice Fellows](#) Recent Graduates Pathways Program could spur research on the important topics of race, gender, and other bias in AI as it relates to law enforcement and justice system applications.

Goal 2: Ensure AI Leadership by Building Partnerships for an AI-Enabled Future

Leveraging AI effectively means changing how Americans work and make policy. Creating cutting-edge technologies does not guarantee that the U.S. government will adopt those technologies effectively and/or regulate them safely. Effective leadership in AI requires not only strategies to develop an AI workforce, but also better partners and partnerships—whether with government, academia, or industry—as well as a national dialogue surrounding AI norms and ethics. The United States needs to **prioritize creating avenues for cross-sector and intergovernmental partnerships. In particular, it needs to bridge the gap between industry**, which has availability to far greater data and resources, and academia, which is critical to basic and translational research.

g) Develop public-private research partnerships to bridge the academia-industry divide.

A critical first step in creating a conducive environment for innovation in AI is bridging the divide between academia and industry. At a time when industry and government are struggling to collaborate in some areas, the academy-industry partnership can help engender trust and make real progress.

At the same time, there are tensions between campuses and companies. While the flow of academics into the private sector may be beneficial in the short run, losing academics in academia will work to the detriment of longer-term, fundamental research. Bringing academia and industry closer together, therefore, will help strike a necessary balance between long-term and short-term research. In order to reform institutions and policies for AI development, multiple stakeholders—including academia and the private sector—need to have a voice.

There is historical precedent for such a move. The United States formed the National Aeronautics and Space Administration (NASA) in mid-1958, after the shock of *Sputnik* in 1957, to create a centralized government organization to ensure American leadership in civilian space exploration. As early as 1958, private companies were involved through contracting while NASA retained control of innovation policy. Over time, however, space policy saw increasing involvement from the private sector and commercialization, such as with the establishment of the “NASA Guidelines Regarding Early Usage of Space for Industrial Purposes,” the NASA Office of Commercial Programs, the Commercial Space Launch Act, and eventually, the formation of the Space Act Agreements. Ultimately of 400,000 people involved in the race to the moon, over 350,000 were employed in the private sector. The United States would not have won the space race had it not been for the flexibility of NASA to involve and leverage the comparative advantages of the private sector. Emerging similarly victorious in the race for AI means bringing all of America’s expertise to bear on AI policy and research.

Penn Input to National Security Commission on AI

Recommendation: Actively facilitate and incentivize cooperation and collaboration on AI—with a focus on data access and compute resources—between academia and the private sector. Providing access to U.S. government data or giving tax benefits to private companies, for example, could incentivize private companies to share their data and compute capabilities with academics and research institutions.

Recommendation: Accelerate efforts to make U.S. government datasets and compute resources available to university researchers.

Penn Input to National Security Commission on AI

h) Replicate successful interagency research models.

In January 2000, President Bill Clinton launched the National Nanotechnology Initiative (NNI) to coordinate the nanotechnology-related activities of 25 different federal departments and independent agencies, managed within the framework of the National Science and Technology Council, under the Office of Science and Technology Policy (OSTP). Prior to the launch of the NNI and as early as 1996, federal agencies such as the NSF and Defense Advanced Research Projects Agency (DARPA) had already begun to launch their own independent programs in nanoscale science and engineering. By 1998, the importance of nanoscience and engineering as a widely applicable, enabling technology (much like AI, it was [likened](#) to the steam engine, transistor, and internet) was becoming more apparent, and a technology evaluation that same year [concluded that](#) “global leadership in nanotechnology was up for grabs.”

Thus, the NNI was born, driven by the [principles](#) of enabling exploration of a new domain of scientific knowledge and incorporating this transformational general-purpose technology into the national technological infrastructure. This initiative developed a national network of research infrastructure in labs to reduce redundancy and to provide access to state-of-the-art infrastructure for the national research enterprise. It motivated the development of new courses, new certifications, and new degree programs across the country and yielded an approach that supports workforce development.

The NNI supports basic and fundamental research initiatives [that are](#) “high risk, high payoff, and broadly enabling,” and guides the [development of](#) “a balanced infrastructure, novel approaches to the education and training of future nanotechnology workers, the ethical, legal, and social implications of nanotechnology, and rapid transfer of knowledge and technology gained from the research and development efforts.”

Today, there is a strong [consensus](#) that the NNI has succeeded in leveraging investments by government, industry, and scientific communities in nanotechnology, significantly contributing to the generation of world-leading research and development as well as to the advancement of knowledge and innovation in the United States. The initiative [spurred](#) a “thriving interdisciplinary nanotechnology community of about 150,000 contributors,” as well as a “flexible R&D infrastructure consisting of about 100 large nanotechnology-oriented R&D centers, networks, and user facilities,” and an “expanding industrial base of about 3,000 companies producing nanotechnology-enabled products.”

Although the initiative is solely a federal interagency initiative, it has served [as a catalyst](#) for America’s broader nanotechnology infrastructure and has led to increased participation by academic, industry, business, civic, governmental, and nongovernmental organizations—all of which operate in a complementary and mutually supportive manner to the direct work of the NNI.

Penn Input to National Security Commission on AI

Despite the overall success of the NNI, the government could have done more and faster. After the initial financial investment, funds for the initiative depended on preexisting agency efforts on nanotechnology. If faster success had been necessary, more new resources would have been essential.

Given that the current status of AI development in the United States is very similar to that of nanotechnology in 1998, an interagency national AI initiative modeled after the NNI and bureaucratically located at a high level, such as under the National Science and Technology Council and OSTP, would prove just as essential for stimulating AI innovation as it did for nanotechnology. The very same rationale that led to the institution of the NNI—an emerging, enabling technology was poised to have dramatic global impact, and the United States wanted to solidify its position as a world leader, but only ad hoc projects at independent agencies existed at the time—currently is up for debate with respect to AI. One limitation of the NNI, however, was its budget model, which slowed the time it took for progress to aggregate. A different approach to investment would be necessary to achieve global leadership in AI in a reasonable time frame.

Furthermore, it would serve the United States well to connect and support its existing projects, so as to not create parallel processes and instead accelerate current successes and aid in filling and shrinking current gaps and challenges. An interagency effort can help to expand and improve such efforts by facilitating collaboration. Currently existing areas of work and projects that such an initiative could encompass include: the [Department of Defense's Joint Artificial Intelligence Center's AI-centric National Mission Initiatives](#), the [Defense Innovation Board's AI Principles](#), the [OSTP's AI research and development programs](#), the [Department of Defense's DARPA AI Basic Research program](#), the [OSTP's STEM Education program](#), and the Department of Commerce's [Commerce Export Controls for AI](#).

Recommendation: Consider launching an interagency AI initiative modeled after the NNI. Ensure that it receives new funding beyond existing initiatives.

Penn Input to National Security Commission on AI

i) *Make long-term investments in people, not just projects.*

When it comes to AI (and STEM research in general), there appears to be a consensus that **there is a [significant lack of investment in basic scientific research](#)**—a necessary component for true innovation—which has put the United States at a potential strategic disadvantage. The lopsidedness in funding, in part, is due to the current way the federal government funds scientific research. As it stands, the current model for research funding is majority project-based and application-driven; it [eschews](#) basic, “blue-sky research”—scientific research in domains where real-world applications are not immediately apparent—in favor of short-term (two- to three-year), low-risk, highly specific projects. Researchers often criticize this current model as being highly [competitive](#) and relatively arbitrary, [wasting time and energy](#), hindering creativity, and impoverishing fundamental or long-term research.

One approach to generating scientific breakthroughs by investing in people comes from the Howard Hughes Medical Institute (HHMI) model. Rather than a fixed end goal, **HHMI prioritizes human capital**. The organization generously funds over long periods of time talented scientists working within a university and allows them to change research direction as the researcher sees fit.

To give a sense of the scope and success of the program, currently HHMI supports nearly 300 “investigators”; in May 2018, HHMI invested \$200 million in just 19 individuals. To date, 29 either current or former HHMI-affiliated investigators have received a Nobel Prize.

The organization practices [risk forgiveness](#), whereas the federal system of recurring grant review is traditionally risk-averse. As explained by [Roberta B. Ness](#), “the HHMI philosophy is to allow investigators to follow ideas through, even if doing so takes a long time. HHMI funds the investigator for many years and appreciates that breakthroughs require hitting walls and making mid-course corrections.” In other words, “The HHMI message to the investigator is: We trust you to be outstanding and internally motivated. The federal message to the investigator is: We provide external motivation at the price that you must constantly prove that we can trust you to produce.”

In a study comparing HHMI- and NIH-sponsored investigators, HHMI scientists changed their research direction far more often than their NIH counterparts—and they produced far more high-impact papers. For [example](#), the United Kingdom, Denmark, Australia, and New Zealand have all employed a hybrid funding model that has a competitive grant program for project-specific funding, while providing stable, multi-year funding for salaries and basic infrastructure dependent on past performance and a history of demonstrated success.

Penn Input to National Security Commission on AI

The federal government should identify people, not projects, in whom they can make long-term investments. While not all funding has to be high-risk and person-centric, introducing mechanisms that would allow the federal government to place big bets and move away from project-only funding would aid greatly in the development of translational AI research.

Recommendation: Create a funding vehicle for artificial intelligence that invests in individual lab leaders for long time periods rather than awarding project-specific grants.

Penn Input to National Security Commission on AI

j) Ensure that AI expertise exists at senior levels of government.

Given that AI leadership will not just require effective research and substantial workforce training but adoption, the **United States needs to ensure that AI leadership has visibility at senior levels of government.** Ensuring that key institutions, committees, and organizations, as well as senior leader staff teams include at least one designated AI expert, would serve as a good way to ensure that AI knowledge is available and that AI implications receive due weight in policy, from foreign to economic.

Moreover, such AI leadership would also support other policymakers who may not have technical expertise by arming them with on-demand AI knowledge as well as familiarizing them with AI-relevant concerns.

Although there are many ways to facilitate this expertise, there is a role for Penn to do what it does best: educate. Along with partners in government and think tanks, Penn scholars could organize a one-week AI boot camp in Washington or Philadelphia for designated AI experts from Congress, the White House, the agencies, and even the Supreme Court. Attendees could earn a certificate after demonstrating knowledge of technical details, policy questions, and ethical issues.

Having an AI-designated member who would convene and filter AI-relevant developments, expertise, and research to bring them to bear on broader policy would be highly effective in streamlining AI policy. Given that AI is an extremely wide-ranging, enabling technology that has the potential to drastically impact nearly every sector of life and a wide array of diverse stakeholders, having AI experts sit in on a range of governing bodies in which all considerations and concerns can be voiced would prove invaluable. The inclusion of such individuals and expertise could strengthen AI policymaking by coordinating, consolidating, and ensuring that AI policy decisions and programs are consistent with America's ultimate strategic goals.

Recommendation: Create AI boot camp for government officials and staffers.

Penn Input to National Security Commission on AI

k) Work with universities and civil society to foster a national dialogue on AI ethics.

Fundamental social norms and practices—such as privacy, fairness, safety, security, and transparency—must inform AI and data science. The rapid emergence of self-driving vehicles and autonomous robots shines the spotlight on important questions of trust and safety.

The private sector has funded the majority of AI research, with most research investment going into AI capability rather than safety. A lack of research into AI safety may create trust issues as well as raise ethical questions about the appropriate use of AI systems. The problems presented by AI are unique due to AI's ability to make independent judgements about situations and operational actions. In order to address these problems, the nation needs to understand and answer questions about appropriate and ethical uses of AI. Ultimately, the government should be able to articulate the principles of use and establish clear guidelines for responsible usage. In order to do so, the government must remain transparent and engage the American public. The federal government should engage academic and other public institutions in a national dialogue on AI ethics, as universities and civil-society organizations are by nature well-poised to host these conversations and are historically collaborative centers of norm creation.

An appropriate analogy is the emergence of the field of bioethics, which started as a philosophical debate among academics and activists in the 1960s and 1970s in response to significant advancement in the medical sciences as well as human experiments, such as those conducted during World War II. [Bioethics](#) originated as a collaborative enterprise in the United States largely in order to differentiate it from medical self-regulation, with participants ranging from philosophers, theologians, lawyers, sociologists, and clinicians to biomedical scientists. Eventually, the debate led to a direct policy outcome, via the National Research Act of 1974 and the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, which eventually produced the 1976 [Belmont Report](#).

Recommendation: Create a board of nonpartisan AI experts to make recommendations on the ethical applications and consequences of AI and to provide oversight for policymaking, assessing the strategy used for technology development and deployment, and evaluating whether technology companies are fair, accountable, transparent, and ethical.