

September 22, 2023

Senator Bill Cassidy, MD
Ranking Member
Senate Committee on Health, Education, Labor, & Pensions

RE: Federation of American Scientists' Comment on "Exploring Congress' Framework for the Future of AI"; Senate Committee on Health, Education, Labor, & Pensions.

The Federation of American Scientists (FAS) is a catalytic, non-partisan, and nonprofit organization committed to using science and technology to benefit humanity by delivering on the promise of equitable and impactful policy. FAS believes that society benefits from a federal government that harnesses science, technology, and innovation to meet ambitious policy goals and deliver impact to the public. We are writing today to provide a response to the Ranking Member of the Senate Committee on Health, Education, Labor, & Pensions (HELP)'s request for information regarding AI in our healthcare system, in the classroom, and in the workplace.¹

Executive Summary

This response provides recommendations on leveraging AI to improve education, healthcare, and the future of work. Key points include:

- In education, the federal government should fund more AI research and establish initiatives like the National Center for Advanced Development in Education. Diverse teams and community-driven design will help ensure that AI use promotes equity.
- For healthcare, Congress should update regulations to enable safe innovation and establish AI Centers of Excellence, starting with maternal health. These can coordinate data sharing and develop AI tools to tackle disparities.

¹ *Exploring Congress' Framework for the Future of AI* (HELP Committee GOP Final AI White Paper), www.help.senate.gov/imo/media/doc/help_committee_gop_final_ai_white_paper1.pdf.

- Regarding labor, AI creates opportunities for job growth and workforce transition. The NIST AI Risk Framework enables responsible adoption. Industry can self-regulate using shared standards to ensure worker protections.

Overall, with thoughtful oversight and human-centric design, AI promises immense benefits across these sectors. But responsible governance is crucial, as is inclusive development and ongoing risk assessment. By bringing together stakeholders, the U.S. can lead in advancing ethical, high-impact applications of AI.

Education

The Federation of American Scientists (FAS) co-leads the Alliance for Learning Innovation (ALI), a coalition of cross-sector organizations seeking to build a stronger, more competitive research and development (R&D) infrastructure in U.S. education. As was noted in the ALI Coalition’s response to White House Office of Science & Technology Policy’s “Request for Information: National Priorities for Artificial Intelligence,”² FAS sees great promise and opportunity for artificial intelligence to improve education, equity, economic opportunity, and national security. In order to realize this opportunity and mitigate risks, we must ensure that the US has a robust, inclusive, and updated education R&D ecosystem that crosscuts federal agencies. Recent breakthroughs in AI, namely large language models like ChatGPT, present new opportunities for teaching and learning. AI has the capacity to grade assessments, give task reminders, provide tailored feedback, draft notes to students’ families, and offer tutoring services. This emerging technology has the potential to lower teacher burnout, increase family engagement, and boost student outcomes.

What Should The Federal Role Be In Supporting AI In Education?

To build a workforce prepared to lead – or at least keep pace with other countries – investments must be made in K-12 educational settings. By infusing

² Alliance for Learning Innovation (7 July 2023). “Response to OSTP Request for Information: National Priorities for Artificial Intelligence.”

https://www.alicoalition.org/_files/ugd/a11dea_288c77acc97b4d299af8528cabcd939c.pdf.

opportunities and funding at the K-12 level, the federal government can be a leader in preparing today's students for tomorrow's challenges.

Research And Development

Specifically, the US government should prioritize funding and supporting R&D in the field of AI to ensure that the U.S. is on the cutting edge of this technology. According to the Stanford Institute for Human-Centered Artificial Intelligence's 2023 AI Index Report, China is the global leader in AI research.³ To keep pace with or surpass China in AI research productivity, we must invest in domestic AI research.

One strong existing Federal example are the AI Institutes supported by the National Science Foundation (NSF) and the U.S. Department of Education (ED).⁴ Earlier this year, NSF and the Institute of Education Sciences (IES) established the AI Institute for Exceptional Children⁵, which capitalizes on the latest AI research to serve children with speech and language pathology needs. Communities would benefit from additional AI Institutes that meet the moment and deliver solutions for today's teaching and learning challenges. Privacy and bias considerations need to be inherent in the research questions being tackled by these and other federally funded AI research efforts.

Rotator Programs

There is also a lot to be learned⁶ from the NSF's rotator program⁶, which is instrumental in ensuring the agency's programs reflect creative ideas from the field, the most advanced methods and diverse teams. More agencies should support rotator programs to ensure they are bringing in diverse and leading talent to address the potential implications and unique needs of diverse communities. These programs could leverage authority from the

³ Maslej, N., et. al (April 2023). *The AI Index 2023 Annual Report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf.

⁴ "National Artificial Intelligence Research Institutes." *NSF - National Science Foundation*, 31 July 2023, <https://new.nsf.gov/funding/opportunities/national-artificial-intelligence-research>.

⁵ "NSF AI Institute for Transforming Education for Children with Speech and Language Processing Challenges." *University at Buffalo*, 27 July 2023, www.buffalo.edu/ai4exceptionaled.html.

⁶ "Rotator Programs." *NSF - National Science Foundation*, <https://new.nsf.gov/careers/rotator-programs>.

Intergovernmental Personnel Act to engage advanced scientific and technical expertise, and build internal capacity.

Expanding Research Grant Programs

Federal agencies, and specifically IES, should build upon the training programs it has for broadening participation and create specific research grant programs for historically Black colleges & universities (HBCUs), tribal colleges & universities (TCUs), and other minority-serving institutions (MSIs) with an emphasis on AI research. While the IES Pathways program⁷ has had success in diversifying education research training programs, more needs to be done at the predoctoral and postdoctoral level.

National Center For Advanced Development In Education

Another key opportunity to support transformational AI research and development in the United States is to establish a National Center for Advanced Development in Education (NCADE).⁸ Modeled after the Defense Advanced Research Projects Agency (DARPA), NCADE would support large-scale, innovative projects that require a more nimble and responsive program management approach than currently in place. Specifically, NCADE would fund projects developed by universities, nonprofits, industry, and innovative organizations, selected based on their potential to create dramatic breakthroughs in learning and teaching, especially for the most underserved populations. Like DARPA, NCADE would be oriented toward ambitious ideas across the academic, public, and private sectors. It would build on bipartisan interest in expanding education R&D, and put outcomes for students at the center of its work. The Center would focus on breakthrough technologies, new pedagogical approaches, innovative learning models, and more efficient, reliable, and valid forms of assessments. This federal investment would spark innovation at state and local levels, and help build local infrastructure to sustain the work. By creating NCADE, Congress can seed the development and use of artificial intelligence to support teaching, personalize learning, support ELL students, and analyze speech and reading.

⁷ *Pathways to the Education Sciences Research Training Program.*

<https://ies.ed.gov/ncer/projects/program.asp?ProgID=95>.

⁸ *IES Director: Innovation in the Education Sciences (the New IES).*

<https://ies.ed.gov/director/remarks/02-02-2023.asp>.

How Is Data That Is Collected During The Use Of These Programs In Schools Used By The AI?

Educational data should be used ethically and responsibly to support student learning. AI systems can help synthesize data to identify needs, but humans must drive decision-making. For example, an AI tutoring system can analyze a student's work to determine knowledge gaps and suggest personalized activities. But the teacher should validate insights and decide on interventions.

At the school district level, AI tools might integrate attendance records, test scores, grades, and other data sources to flag at-risk students. Administrators would then investigate causes and connect students with supports like counseling. Importantly, AI should never autonomously determine placements or discipline. Humans must be actively involved.

Any use of student data by AI must be transparent and give parents visibility. Strict data governance policies should be in place, limiting access and handling. Data collection should be opt-in only and not used to profile students. Overall, AI should complement human judgment - providing insights but not replacing the nuance of teachers and administrators acting in students' best interests. With thoughtful design and oversight, data can inform teaching while protecting privacy.

What Are The Best Practices Currently Being Used To Ensure That AI Systems Are Designed, Developed, And Deployed In A Manner That Protects People's Rights And Safety?

First and foremost, we need to ensure that underserved communities, minors, individuals with disabilities and the civil rights organizations that support them are at the table throughout the design process for AI tools and products. In particular, we need to ensure that research is led and driven locally and by those who are closest to the challenges, namely educators, parents, students, and local and state leaders.

Diverse teams produce better tools and technologies and lead to greater innovations and breakthroughs. Microsoft, in its Inclusive Design Toolkit,⁹ refers to this as “learning from diversity,” where including more people in the design process and widening the talent pool is key to addressing important societal challenges. As our society moves to quickly adopt artificial intelligence and machine learning in ways that impact all areas of our lives, there is both potential for greater inclusion or exclusion depending on the way the AI is designed and implemented. Including people with disabilities and other differences in the design of the AI may help safeguard against the negative impacts of bias in AI, which is essential if we are to build an AI-driven future that is equitable for everyone in society.

When thoughtfully and inclusively designed, AI has the potential to enhance equity by providing more personalized learning for students and by supporting educators to address the individual and diverse needs in their classrooms. For example, AI could be utilized in teacher preparation programs to ensure that educators have access to more diverse experiences during their pre-service experiences. AI can also provide benefits and services to students and families who currently do not have access to those resources due to a lack of human capital.

Attention should be paid to the design, evaluation and implementation process consistent with the recommendations in the recent report from the Office of Education Technology, “Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations.”¹⁰ AI works best when it is addressing a particular need in teaching and learning or a particular resource gap.

Labor

AI will transform the future of work, creating new jobs even as it automates some roles. Thoughtful governance can maximize opportunities while mitigating risks. Industry collaboration and AI standards will enable responsible adoption.

⁹ Microsoft Inclusive Design. <https://inclusive.microsoft.design>.

¹⁰ Cardona, M. “Artificial Intelligence and the Future of Teaching and Learning.” *Office of Educational Technology*, U.S. Department of Education, <https://www2.ed.gov/documents/ai-report/ai-report.pdf>.

With proactive policymaking, automation can augment human potential and productivity.

What Role Will AI Play In Creating New Jobs?

AI can serve as a powerful tool for workforce systems, employers, and employees alike in order to drive job creation and upskilling. For instance, investment in large language learning models that scrape and synthesize real-time labor market information (LMI) can be used to better inform employers and industry consortia about pervasive skills gaps. This information could be used by workers to make better informed training and education decisions, allowing for more efficient allocation of labor resources. Currently, most advanced real-time LMI products exist behind paywalls, but Congress should consider investing in the power of this information as a public good to forge a more competitive labor market.

The wide-scale commercialization of AI/ML-based products and services will also create new types of jobs and occupations for workers. Contrary to popular belief, many industries that face some level of automation will still require trained employees to pivot to emerging needs in a way that offsets the obsolescence of other roles. For example, the trucking industry is facing a labor crisis, with a turnover rate nearing 100 percent.¹¹ With this challenge comes an opportunity for autonomous vehicles to play a bigger role in the trucking industry, which requires different skills. To meet this need, TuSimple, the first self-driving tech company to go public in the United States, approached Pima Community College in Arizona, to create the first autonomous driving certificate for truck drivers.¹² Through place-based partnerships between employers and training institutions (e.g., community colleges, work-based learning programs, etc.), localities can reinvest in their workers to provide transition opportunities and close labor market gaps.

¹¹ McNally, S. "Turnover Rate at Large Truckload Carriers Rises in First Quarter." *American Trucking Associations*, 6 June 2018,

www.trucking.org/news-insights/turnover-rate-large-truckload-carriers-rises-first-quarter.

¹² "Autonomous Vehicle Driver & Operations Specialist." *Pima Community College*, <https://www.pima.edu/academics-programs/degrees-certificates/industry-manuf-construction/truck-driving/autonomous-vehicle-specialist-cert/index.html>.

What role will AI standards, such as the National Institute of Standards and Technology AI Risk Management Framework, play in regulatory and self-regulatory efforts?

AI standards will serve as a crucial foundation as the US government and industries navigate AI's impacts on the workforce. The NIST AI Risk Management Framework¹³ provides a methodology for organizations to assess and mitigate risks across the AI lifecycle. This could enable more responsible automation in HR contexts - for example, helping ensure bias mitigation in algorithmic hiring tools. On the policy side, lawmakers drafting regulations around AI and employment will likely reference and even codify elements of the Framework. The recent federal Blueprint for an AI Bill of Rights¹⁴ similarly draws upon existing voluntary AI standards in articulating principles for labor impacts.

Thoughtful frameworks like the NIST AI RMF allow regulators, companies, and workers to better align on ethical deployment of AI. Unions and other worker advocacy groups are already using standards as a basis to demand responsible adoption.¹⁵ Overall, the maturation of AI standards promises more balanced, evidence-based policymaking around automation's labor effects. With proactive collaboration between stakeholders, standards can help ensure technologies boost productivity while respecting worker rights and human dignity.

On the industry side, responsible technology leaders are already using the NIST AI RMF for self-regulation.¹⁶ By proactively auditing and mitigating risks in internal AI systems, companies can build public trust and reduce the need for excessive government intervention. HR departments in particular are wise to adopt such standards in screening, hiring, performance evaluation, and other workforce AI tools. Though policymakers still have an oversight role, widespread

¹³ (26 January 2023) *AI Risk Management Framework*. National Institute of Standards and Technology. <https://www.nist.gov/itl/ai-risk-management-framework>.

¹⁴ "Blueprint for an AI Bill of Rights." *The White House*, The United States Government, 16 Mar. 2023, www.whitehouse.gov/ostp/ai-bill-of-rights/.

¹⁵ "A Worker-Centered Digital Trade Agenda." *AFL-CIO*, 7 Feb. 2023, <https://aflcio.org/worker-centered-digital-agenda>.

¹⁶ "Governing AI: A Blueprint for the Future." *Microsoft*, 25 May 2023, <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RW14Gtw>.

self-regulation using shared frameworks is at this point the most efficient path for safe and responsible AI across the labor market.

Health Care

Artificial intelligence has a long history in the life sciences. By the turn of the century, bioinformaticians (data scientists for biological data) were already using AI in genome analysis.¹⁷ One focus of AI tools for biology has been on proteins. Predicting the shape of a protein has long been a critical challenge. In 2020, Alphabet's DeepMind published AlphaFold 2¹⁸ as an AI-enabled software tool capable of doing just that. While not perfect, scientists have been able to use it and related tools to predict the shape of proteins faster and even to create new proteins optimized for specific applications.

Of course, the applications of AI in biotechnology extends beyond proteins. Medical researchers have taken advantage of AI to identify new biomarkers and leverage AI to improve diagnostic tests. Industrial biotechnology researchers are exploring the use of AI to optimize biomanufacturing processes to improve yield. In other natural sciences, AI can even drive entire courses of experiments with minimal human input,^{19,20} with biological labs in development. Unfortunately, these same tools and capabilities could also be misused to cause harm by actors trying to develop toxins, pathogens, and other potential bio risks.

¹⁷ Zhang, M. Q. (2000). Discriminant analysis and its application in DNA sequence motif recognition. *Briefings in bioinformatics*, 1(4), 331-342.

<https://academic.oup.com/bib/article/1/4/331/2530002>

¹⁸ Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., ... & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583-589.

<https://www.nature.com/articles/s41586-021-03819-2>

¹⁹ Carnegie Mellon University and Emerald Cloud Lab to Build World's First University Cloud Lab. (2021). Carnegie Mellon University.

<https://www.cmu.edu/news/stories/archives/2021/august/first-academic-cloud-lab.html>

²⁰ Boiko, D. A., MacKnight, R., & Gomes, G. (2023). Emergent autonomous scientific research capabilities of large language models. *arXiv preprint arXiv:2304.05332*.

What Updates To The Regulatory Frameworks For Drugs And Biologics Should Congress Consider To Facilitate Innovation In AI Applications?

Congress has an opportunity to update regulations to enable responsible innovation and oversight for AI applications in biopharma. For example, the FDA could issue detailed guidance clarifying pathways for developers to validate and seek approval for AI-designed drugs, biologics, or medical devices through traditional clinical trials. The guidance could specify validation standards to ensure AI-designed products are safe and effective. Congress could also consider expanding the FDA's mandate and capacity to require upfront risk assessments before deployment of particularly high-risk or dual-use bio-AI systems.²¹ This approach is currently used by DARPA for some autonomous and biological technologies.

Additionally, thoughtful reporting requirements could be instituted for entities developing advanced bio-AI models above a certain capability threshold. This transparency would allow for monitoring of dual-use risks while avoiding overregulation of basic research. Congress could also empower the Biomedical Advanced Research and Development Authority (BARDA) to establish a public-private partnership specifically focused on biosecurity for AI in the life sciences. With prudent oversight mechanisms in place, data sharing and scientific collaboration could continue unimpeded to drive progress in biomedicine.

How Can The FDA Improve The Use Of AI In Medical Devices? What Updates To The Regulatory Frameworks For Medical Devices Should Congress Consider To Facilitate Innovation In AI Applications While Also Ensuring That Products Are Safe And Effective For Patients?

AI holds promise to improve health outcomes, but a number of oversight measures will help to ensure safety and efficacy. The FDA should issue guidance

²¹ Letter by Representative Anna G. Eshoo to National Security Advisor Jake Sullivan and Director of White House Office of Science and Technology Policy Arati Prabhakar. 2022.
<https://eshoo.house.gov/media/press-releases/eshoo-urges-nsa-ostp-address-biosecurity-risks-caused-ai>

on subpopulation performance testing and real-world monitoring of AI tools, and Congress should clarify liability for AI device errors. With thoughtful regulation, AI in medicine can augment human expertise while protecting patients.

Ensuring That Analysis Of Subpopulation Performance Is A Key Component Of The Review Process For AI Tools

Analyzing data on the subpopulation performance of medical devices should be one key component of any comprehensive effort to advance equity in medical innovation.²² We appreciate the recommendations in the GOP HELP white paper asking developers to document the performance of their devices on various subpopulations when considering updates and modifications. It will be essential to assess subpopulation performance to mitigate harms that may otherwise arise—especially if an argument for equity is made for a certain product.

More Specific Guidance On Items Of Regional Concern

The GOP HELP white paper usefully addresses some instances in which AI systems can be adapted to the conditions in which they are deployed. This is an important issue, as many systems can vary substantially in efficacy depending on the local contexts in which they are used. To provide one concrete example of this issue from a recent study, “Artificial intelligence (AI) performance in breast cancer screening was affected by mammography equipment and software used, highlighting the importance of local clinical settings and technology for effective AI implementation.”²³

If devices are being marketed as broadly effective without explanation or attention being given to the ways in which these systems may fall short in contexts that are notably different from the training environment, devices may be sold while falling far short of their promise and thus could lead to significant harm. To ensure safety and avoid harms of this type, special attention and more

²² Wickerson, G. (2022). Combating Bias In Medical Innovation. Federation of American Scientists. <https://fas.org/publication/combating-bias-in-medical-innovation/>

²³ de Vries, C. F., Colosimo, S. J., Staff, R. T., Dymiter, J. A., Yearsley, J., Dinneen, D., ... & iCAIRD Radiology Collaboration. (2023). Impact of Different Mammography Systems on Artificial Intelligence Performance in Breast Cancer Screening. *Radiology: Artificial Intelligence*, 5(3), e220146. <https://pubs.rsna.org/doi/10.1148/ryai.220146>

specific guidance on how such harms can be avoided must be given. If errors are being noticed by clinicians, such as those working in clinics with older equipment, there should be a process in place to trigger post-market review and revisions to the tool.

Clarifying The Role Of Real-World Evidence In Approvals

Locating concerns regarding performance in subpopulations and within different medical environments will most likely involve the collection of real-world evidence regarding the performance of these tools in the wild. The role of real-world evidence²⁴ in the regulatory approval process for market surveillance and updates should be defined more clearly in this guidance. Is there a preference for real-world evidence over other forms of data? The GOP HELP white paper includes a call for algorithmic impact assessments, which could provide a helpful framework to monitor algorithm performance and determine the need for updates in a proactive, manufacturer-driven model. Moreover, there should be a consistent flow of real-world evidence for AI devices in actual use scenarios, not just during adverse events.

Clarifying Legal Liability For Errors In AI Tools

Another issue that remains unresolved is establishing a liability regime for errors made by AI tools.^{25,26} The GOP HELP white paper outlines an approach for how manufacturers can choose to modify their devices but does not offer any explanation of conditions under which the FDA may require a change or an update to be made to a device. For example, if an independent study suggested that a device was erring in a way that was leading to harm, could the FDA require that the manufacturer institute changes outlined in the Predetermined Change Control Plan?

²⁴ (2023). U.S. Food and Drug Administration. Real-World Evidence. <https://www.fda.gov/science-research/science-and-research-special-topics/real-world-evidence>.

²⁵ Khullar, D., Casalino, L. P., Qian, Y., Lu, Y., Chang, E., & Aneja, S. (2021). Public vs physician views of liability for artificial intelligence in health care. *Journal of the American Medical Informatics Association*, 28(7), 1574-1577. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8279784/>

²⁶ Maliha, G., & Parikh, R. B. (2022). Who Is Liable When AI Kills. *Scientific American*. <https://www.scientificamerican.com/article/who-is-liable-when-ai-kills/>

How Can AI Be Best Adopted To Not Inappropriately Deny Patients Care?

AI Centers of Excellence could be established to develop demonstration AI tools for specific care populations and care environments. For example, FAS has published a Day One Memo proposing an AI Center of Excellence for Maternal Health²⁷ to bring together data sources, then analyze, diagnose, and address maternal health disparities, all while demonstrating trustworthy and responsible AI principles. The benefit of AI Centers of Excellence are two-fold: they provide an opportunity for coordination across the federal government, such as for federal data sources held by Qualified Health Information Network (QHIN) and federal data from the CDC, Centers for Medicare and Medicaid (CMS), Office of Personnel Management (OPM), Healthcare Resources and Services Agency (HRSA), NIH, United States Department of Agriculture (USDA), Housing and Urban Development (HUD), the Veterans Health Administration, and Environmental Protection Agency (EPA); and they can evaluate existing datasets to establish high-priority, high-impact applications of AI-enabled research for improving clinical care guidelines and tools for healthcare providers.

The following examples demonstrate how AI might help address maternal health disparities, based on priority areas informed by clinicians in the field:²⁸

- AI implementation should be explored for analysis of electronic health records from the VHA and QHIN to predict patients who have a higher risk of pregnancy and/or delivery complications.
- Drawing on the robust data collection and patient surveillance capabilities of the VHA and HRSA, AI should be explored for the deployment of digital tools to help monitor patients during pregnancy to ensure adequate and consistent use of prenatal care.

²⁷ Sennaar, K. & Wickerson, G. (2023). Establishing An AI Center Of Excellence To Address Maternal Health Disparities. Federation of American Scientists.
<https://fas.org/publication/establishing-an-ai-center-of-excellence-to-address-maternal-health-disparities/>

²⁸ Padilla, C., Abir, G., Zakowski, M., & Carvalho, B. (2021). How AI Could Help Doctors Reduce Maternal Mortality. Harvard Business Review.
<https://hbr.org/2021/08/how-ai-could-help-doctors-reduce-maternal-mortality>

- Using VHA data and QHIN data, AI should be explored in supporting patient monitoring in instances of patient referrals and/or transfers to hospitals that are appropriately equipped to serve high-risk patients, following guidelines provided by the American College of Obstetricians and Gynecologists.
- Data on housing from HUD, rural development from the USDA, environmental health from the EPA, and social determinants of health research from the CDC should be connected to risk factors for maternal mortality in the academic literature to create an AI-powered risk algorithm.
- Understand the power of payment models operated by CMS and OPM for novel strategies to enhance maternal health outcomes and reduce maternal deaths.

The AI Center of Excellence model demonstrates the power of coordinating and thoughtfully applying AI tools across disparate federal data sources to address urgent public health needs. Similar centers could be established to tackle other complex challenges at the intersection of health, environmental, socioeconomic, and demographic factors and to provide better, more equitable care. For example, an AI Center focused on childhood asthma could integrate housing data, EPA air quality data, Medicaid records, and school absenteeism data to understand and predict asthma triggers. The key is a laser focus on high-impact AI applications, grounded in the real needs of vulnerable communities and guided by domain experts.

Conclusion: Harnessing the Promise of AI

Artificial intelligence holds tremendous potential to transform education, healthcare, and work for the better. But realizing these benefits in an equitable, ethical way requires proactive collaboration between policymakers, researchers, civil society groups, and industry.

The recommendations outlined here aim to strike a balance - enabling innovation and growth while centering human needs and mitigating risks. This requires robust funding for R&D, modernized regulations, voluntary standards,

and inclusive design principles. Ongoing oversight and impact evaluation will be key, as will coordination across agencies and stakeholders.

With wise governance and a commitment to advancing AI for social good, the U.S. can lead in developing cutting-edge yet trustworthy technologies. By keeping people at the heart of system design and policymaking, we can build an AI-powered future that promotes opportunity, equity, and human potential. The ideas and frameworks in this response provide a roadmap toward that goal.

We at FAS would welcome the chance to discuss these ideas further with the Senate HELP Committee Minority. Our organization stands ready to serve as a resource on AI policy issues, and to help craft solutions that harness AI's immense potential while upholding our shared values. We invite your office to connect with us to explore these proposals in more depth.

Divyansh Kaushik
Associate Director, Emerging Technologies
and National Security
dkaushik@fas.org

Jack Titus
Fellow, AI Policy
jtitus@fas.org