



ITI



AI POLICY PRINCIPLES Executive Summary

Artificial Intelligence (AI) is a suite of technologies capable of learning, reasoning, adapting, and performing tasks in ways inspired by the human mind. With access to data and the computational power and human ingenuity required to extract increasing value from it, researchers are building intelligent software and machines to enhance human productivity and empower people everywhere. Startups, medium-sized companies, and larger technology companies have all developed AI systems to help solve some of society's most pressing problems, from medical diagnosis to education to economic productivity and empowerment.

While it is impossible to predict the full transformational nature of AI, like technological evolutions before it, we expect the potential implications to be vast. To ensure that AI can deliver its greatest positive potential, the Information Technology Industry Council (ITI) — the global voice of the tech sector — takes industry's responsibility seriously to be a catalyst for preparing for an AI world. In our Policy Principles, we outline specific areas where industry, governments, and others can collaborate, as well as specific opportunities for public-private partnership. To advance these principles, which we expect will evolve alongside AI technology, we acknowledge the following:

Industry's Responsibility in Promoting Responsible Development and Use: We recognize our responsibility to integrate principles into the design of AI technologies, beyond compliance with existing laws. While the potential benefits to people and society are amazing, AI researchers, subject matter experts, and stakeholders should continue to spend a great deal of time working to ensure the responsible design and deployment of AI systems, including addressing safety and controllability mechanisms, use of robust and representative data, enabling greater interpretability and recognizing that solutions must be tailored to the unique risks presented by the specific context in which a particular system operates.

The Opportunity for Governments to Invest In and Enable the AI Ecosystem: We encourage robust support for research and development (R&D) to foster innovation through incentives and funding. As the primary source of funding for long-term, high-risk research initiatives, we support governments' investment in research fields specific or highly relevant to AI, including: cyber-defense, data analytics, detection of fraudulent transactions or messages, robotics, human augmentation, natural language processing, interfaces, and visualizations. We also encourage governments to evaluate existing policy tools and use caution before adopting new laws, regulations, or taxes that may inadvertently or unnecessarily impede the responsible development and use of AI. This extends to the foundational nature of protecting source code, proprietary algorithms, and other intellectual property. Failure to do so could present a significant cyber risk.

The Opportunity for Public-Private Partnerships (PPPs): Many emerging AI technologies are designed to perform a specific task, assisting human employees and making jobs easier. Our ability to adapt to rapid technological change, however, is critical. That is why we must continue to be prepared to address the implications of AI on the existing and future workforce. By leveraging PPPs – especially between industry partners, academic institutions, and governments – we can expedite AI R&D, democratize access, prioritize diversity and inclusion, and prepare our workforce for the jobs of the future.



ITI



AI POLICY PRINCIPLES

Artificial Intelligence (AI) is a suite of technologies capable of learning, reasoning, adapting, and performing tasks in ways inspired by the human mind. With access to data and the computational power and human ingenuity required to extract increasing value from it, researchers are building intelligent software and machines to enhance human productivity and empower people everywhere.


We are already experiencing how AI **benefits people, society, and the economy** in a diverse array of fields. AI systems assist in medical diagnostics, alerting doctors to early warning signs, and helping personalize patient treatments. They increase accessibility, fueling software programs that make digital content accessible to people with disabilities, such as helping the blind “read” millions of photos or websites on the internet. And intelligent systems already monitor huge volumes of economic transactions – identifying potential fraud in real time and saving consumers millions of dollars.

By pairing the power of AI computing with land cover maps, weather forecasts, and soil data, technology can empower people with the data and tools they need to better conserve lands, improve ecosystems, and increase agricultural yields. AI-powered machines can even make dangerous or difficult tasks safer for people, opening new environments that were previously inaccessible to human exploration.

Startups, medium-sized companies, and larger technology companies have all developed AI systems to help solve some of society’s most pressing problems. By allowing smaller businesses to do more with less, AI will jumpstart small businesses, helping them take risks and grow at faster rates than ever before.

Like other transformative technological evolutions before it, it is impossible to fully predict the impact of AI, but like the development of the internet, we expect the potential implications to be vast. In the United States alone, the market for AI technologies that analyze unstructured data is projected to reach \$40 billion by 2020, potentially generating more than \$60 billion worth of productivity improvements per year. By their very nature, these innovations create new products and services that did not exist before. By 2025, AI technologies are expected to add between \$7.1 trillion and \$13.17 trillion to the global economy.

These transformations should not cloud the fact that AI remains an active area of research that is constantly evolving and improving. As it evolves, we take our responsibility seriously to be a catalyst for preparing for an AI world, including seeking solutions to address potential negative externalities and helping to train the workforce of the future.



To ensure that AI is able to deliver its greatest positive potential, the Information Technology Industry Council (ITI) — representing the technology sector’s leading companies — urges collaboration among stakeholders across public and private sectors. We, as an industry, acknowledge the need to develop dialogues with governments and other interested parties to make this an inclusive process at every stage. Outlined below are specific areas where industry and governments can collaborate, followed by specific opportunities for public-private partnerships (PPPs). To advance these principles, which we expect will evolve alongside AI technology, we acknowledge the following:

Our Responsibility: Promoting Responsible Development and Use.

Responsible Design and Deployment: We recognize our responsibility to integrate principles into the design of AI technologies, beyond compliance with existing laws. While the potential benefits to people and society are amazing, AI researchers, subject matter experts, and stakeholders should and do spend a great deal of time working to ensure the responsible design and deployment of AI systems. Highly autonomous AI systems must be designed consistent with international conventions that preserve human dignity, rights, and freedoms. As an industry, it is our responsibility to recognize potentials for use and misuse, the implications of such actions, and the responsibility and opportunity to take steps to avoid the reasonably predictable misuse of this technology by committing to ethics by design.



Safety and Controllability: Technologists have a responsibility to ensure the safe design of AI systems. Autonomous AI agents must treat the safety of users and third parties as a paramount concern, and AI technologies should strive to reduce risks to humans. Furthermore, the development of autonomous AI systems must have safeguards to ensure controllability of the AI system by humans, tailored to the specific context in which a particular system operates.

Robust and Representative Data: To promote the responsible use of data and ensure its integrity at every stage, industry has a responsibility to understand the parameters and characteristics of the data, to demonstrate the recognition of potentially harmful bias, and to test for potential bias before and throughout the deployment of AI systems. AI systems need to leverage large datasets, and the availability of robust and representative data for building and improving AI and machine learning systems is of utmost importance.



Interpretability: We are committed to partnering with others across government, private industry, academia, and civil society to find ways to mitigate bias, inequity, and other potential harms in automated decision-making systems. Our approach to finding such solutions should be tailored to the unique risks presented by the specific context in which a particular system operates. In many contexts, we believe tools to enable greater interpretability will play an important role.

Liability of AI Systems Due to Autonomy: The use of AI to make autonomous consequential decisions about people, informed by – but often replacing decisions made by – human-driven bureaucratic processes, has led to concerns about liability. Acknowledging existing legal and regulatory frameworks, we are committed to partnering with relevant stakeholders to inform a reasonable accountability framework for all entities in the context of autonomous systems.

The Opportunity for Governments: Investing and Enabling the AI Ecosystem.

Investment in AI Research and Development: We encourage robust support for research and development (R&D) to foster innovation through incentives and funding. As the primary source of funding for long-term, high-risk research initiatives, we support governments' investment in research fields specific or highly relevant to AI, including: cyber-defense, data analytics, detection of fraudulent transactions or messages, robotics, human augmentation, natural language processing, interfaces, and visualizations.



Flexible Regulatory Approach: We encourage governments to evaluate existing policy tools and use caution before adopting new laws, regulations, or taxes that may inadvertently or unnecessarily impede the responsible development and use of AI. As applications of AI technologies vary widely, overregulating can inadvertently reduce the number of technologies created and offered in the marketplace, particularly by startups and smaller businesses. We encourage policymakers to recognize the importance of sector-specific approaches as needed; one regulatory approach will not fit all AI applications. We stand ready to work with policymakers and regulators to address legitimate concerns where they occur.

Promoting Innovation and the Security of the Internet: We strongly support the protection of the foundation of AI, including source code, proprietary algorithms, and other intellectual property. To this end, we believe governments should avoid requiring companies to transfer or provide access to technology, source code, algorithms, or encryption keys as conditions for doing business. We support the use of all available tools, including trade agreements, to achieve these ends.

Cybersecurity and Privacy: Just like technologies that have come before it, AI depends on strong cybersecurity and privacy provisions. We encourage governments to use strong, globally-accepted and deployed cryptography and other security standards that enable trust and interoperability. We also promote voluntary information-sharing on cyberattacks or hacks to better enable consumer protection. The tech sector incorporates strong security features into our products and services to advance trust, including using published algorithms as our default cryptography approach as they have the greatest trust among global stakeholders, and limiting access to encryption keys. Data and cybersecurity are integral to the success of AI. We believe for AI to flourish, users must trust that their personal and sensitive data is protected and handled appropriately. AI systems should use tools, including anonymized data, de-identification, or aggregation to protect personally identifiable information whenever possible.



Global Standards and Best Practices: We promote the development of global voluntary, industry-led, consensus-based standards and best practices. We encourage international collaboration in such activities to help accelerate adoption, promote competition, and enable the cost-effective introduction of AI technologies.

The Opportunity for Public-Private Partnerships: Promoting Lifespan Education and Diversity.

Democratizing Access and Creating Equality of Opportunity: While AI systems are creating new ways to generate economic value, if the value favors only certain incumbent entities, there is a risk of exacerbating existing wage, income, and wealth gaps. We support diversification and broadening of access to the resources necessary for AI development and use, such as computing resources, education, and training, including opportunities to participate in the development of these technologies.



Science, Technology, Engineering and Math (STEM) Education: Current and future workers need to be prepared with the necessary education and training to help them succeed. We recognize that delivering training is critical and will require significant investment, not only in STEM education, but also in understanding human behavior via the humanities and social sciences. To ensure employability of the workforce of the future, the public and private sectors should work together to design and deliver work-based learning and training systems, and advance approaches that provide students with real work experiences and concrete skills. In conjunction, prioritizing diversity and inclusion in STEM fields, and in the AI community specifically, will be a key part in ensuring AI develops in the most robust way possible.



Workforce: There is concern that AI will result in job change, job loss, and/or worker displacement. While these concerns are understandable, it should be noted that most emerging AI technologies are designed to perform a specific task and assist rather than replace human employees. This type of augmented intelligence means that a portion, but most likely not all, of an employee's job could be replaced or made easier by AI. While the full impact of AI on jobs is not yet fully known, in terms of both jobs created and displaced, an ability to adapt to rapid technological change is critical. We should leverage traditional human-centered resources as well as new career educational models and newly developed AI technologies to assist both the existing workforce and future workforce in successfully navigating career development and job transitions. Additionally, we must PPPs that significantly improve the delivery and effectiveness of lifelong career education and learning, inclusive of workforce adjustment programs. We must also prioritize the availability of job-driven training to meet the scale of need, targeting resources to programs that produce strong results.

Public Private Partnership: PPPs will make AI deployments an attractive investment for both government and private industry, and promote innovation, scalability, and sustainability. By leveraging PPPs – especially between industry partners, academic institutions, and governments – we can expedite AI R&D and prepare our workforce for the jobs of the future.