

Testimony of Dr. Amir Khosrowshahi, Vice President and Chief Technology Officer, Intel Corporation, Artificial Intelligence Products Group, before the U.S. House of Representatives, Committee on Oversight and Government Reform, Subcommittee on Information Technology, Hearing “Game Changers: Artificial Intelligence Part I”

February 14, 2018

Good afternoon, Chairman Hurd, Ranking Member Kelly, and members of the House Committee on Oversight and Government Reform, Subcommittee on Information Technology. My name is Amir Khosrowshahi, and I am the Vice President and Chief Technology Officer of Intel Corporation’s Artificial Intelligence Products Group. Prior to joining Intel, I co-founded Nervana Systems, which created new hardware and software technology to allow companies to build Artificial Intelligence solutions. I have an academic background in neuroscience and machine learning from the University of California, Berkeley, and in mathematics and physics from Harvard University.

We are here today to discuss artificial intelligence. The term, “artificial intelligence” was first coined in the 1950’s and until recently has been only an aspirational concept. Thanks to technological advances such as increased computing capability, large datasets, and innovations in algorithms, AI is now beginning to be a fixture in our daily lives. For instance, speech recognition, recommendation engines, and bank fraud detection systems are all features that utilize AI to make our lives more convenient. As CTO of Intel’s AI Products Group, I can attest that despite this progress, the AI industry is still very much in the early days. However, as AI continues to advance, we must recognize that it will also cause shifts within technology, and within our society.

My work at Intel focuses on employing deep learning hardware chips and software algorithms to make inferences with data sets that can be applied to real-world scenarios offering benefits to organizations. In addition we spend a lot of time thinking about how AI can be harnessed to do good. There are several examples of how deep learning algorithms are currently being deployed to solve some of the world’s most pressing challenges and benefitting society in turn.

AI solutions are poised to revolutionize the tools doctors employ to treat their patients. According to the Centers for Disease Control, about one in four deaths each year in the United States is caused by heart disease.¹ Detecting and treating various types of heart disease is tricky, and since doctors have long had to rely on experience and instinct in making diagnoses, patients were left with varying results.

Take for example the situation of a doctor treating a patient presenting common symptoms of heart failure. Making an accurate diagnosis is critical, and conditions like pericarditis and cardiomyopathy are often difficult to distinguish and therefore difficult to treat accurately. Even experienced doctors make the correct diagnosis only 3 out of 4 times. For doctors with less experience, their rate of correct diagnoses falls to 50 percent – as accurate as flipping a coin. In a recent experiment, researchers using AI spotted the difference between pericarditis and cardiomyopathy nine out of 10 times.² In this regard, artificial intelligence democratizes expert diagnoses for patients and doctors everywhere in the world.

¹ <https://www.cdc.gov/heartdisease/facts.htm> last accessed January 28, 2018

² <https://www.intel.com/content/www/us/en/healthcare-it/article/improved-diagnosis.html> (last accessed January 29, 2018)

AI is also contributing to agriculture improvements, providing tools to farmers to increase crop yield. The world's population will increase to 8.5 billion by the year 2030.³ By 2050, we will need to produce at least 50 percent more food to feed the growing number of people in the world.⁴ Meeting this demand will become increasingly difficult, for as populations climb, the supply of land decreases and societies must feed more people with less land to grow crops. Thankfully, with sensors, drones, robots, and advanced compute power, AI applications provide tools to farmers to deal with this problem. Intel technology helps farmers maximize crop yields, reduce environmental impact, and meet growing demand. "Precision agriculture" links sensor data on the ground to nearby Intel® IoT Gateways. The gateways act on that data locally and then send the collected data to cloud service providers, where they are analyzed using solutions from partners. These results should help farmers improve yields, reduce consumption of resources like water and fertilizer, and improve crop quality.

AI is also enabling law enforcement to combat child abuse. When Internet service providers identify instances of suspicious online activity pointing to child exploitation, they send these reports to the National Center for Missing and Exploited Children (NCMEC) - the nation's clearinghouse and comprehensive reporting center for all issues related to the prevention of and recovery from child victimization. From there, the NCMEC combs through these tips, sending legitimate leads to local law enforcement agencies. The problem is, NCMEC has a team of only 25 analysts employed to go through over 8 million tips. Intel partnered⁵ with NCMEC to deliver an AI solution which helps analysts get the right information to the right jurisdiction. With this technology, NCMEC can scan sites for suspicious content, store massive volumes of data, run a variety of queries, and share the data across the organization's applications. AI helps automate and speed up the process, reducing the typical 30-day turnaround time to handle a report to just a day or two, which can save a child's life.⁶

As technologies evolve, they will undoubtedly continue to have an impact on our daily lives. Some have hailed the development of AI as evidence of the fourth industrial revolution. This analogy to the industrial revolution is telling – at that time through the turn of the nineteenth century, people were concerned with the same societal and economic shifts that are now being related to AI. Back then, some were concerned that they'd lose their jobs while others were hopeful for a 10 hour workweek.⁷ In reality, while some industries – like horse-drawn carriages – faced severe displacement, others, like the automobile industry, flourished from the development of the combustion engine, creating many new

³ <http://www.un.org/sustainabledevelopment/blog/2015/07/un-projects-world-population-to-reach-8-5-billion-by-2030-driven-by-growth-in-developing-countries/> (last accessed January 29, 2018)

⁴ <https://www.intel.com/content/www/us/en/big-data/article/agriculture-harvests-big-data.html> (last accessed January 29, 2018)

⁵ <https://www.intel.com/content/www/us/en/analytics/artificial-intelligence/article/ai-helps-find-kids.html> (last accessed January 29, 2018)

⁶ <https://iq.intel.com/artificial-intelligence-is-good-for-society/> last accessed January 30, 2018

⁷ Statement from Naveen Rao, Intel VP and GM of Intel's Artificial Intelligence Products Group at The Atlantic's event, The Innovation Game, November 7, 2017. Statement at approximately 26:10 of panel. Video available here:

https://www.youtube.com/watch?v=aSDojU2HFWg&list=PLWj46yNDLyTV_V_Z5TrHMv1WE5WvK4QK-&index=2 last accessed February 6, 2018

jobs in the process. What didn't happen, despite that shift, was a 10 hour workweek. Indeed, with widespread improvements in productivity, there also tends to be creation of new jobs as well as development of new categories of jobs. As a society, our goal now should be figuring out how we can harness AI, tapping into the many benefits it holds, while mitigating the impacts of disruptions. And we need to embrace changes in technology and empower people to participate in an AI-fueled economy.

Governments will play a major role if we are to harness the benefits of AI while mitigating possible impacts. Broadly, as with any technological innovation, governments should focus their attention now on enabling the development of AI. We are in the early days of an innovation of a technology that can do tremendous good. Governments should make certain to encourage this innovation, and they should be wary of regulation that will stifle its growth. More specifically, Governments can assist in aiding the further development of this technology, ensuring citizens are equipped to participate in this emerging work economy, and safeguarding against potential disruptions to our way of life.

First, at the Federal level, the United States Government can play an important role in enabling the further development of AI technology. Since data is fuel for AI, the U.S. Government should embrace open data policies. For us to realize AI's benefits, researchers need to have access to large data sets. Some of the most comprehensive data sets are currently owned by the federal government. This data is a taxpayer-funded resource which, if made accessible to the public, could be utilized by researchers to train algorithms which would undergird future AI solutions. The OPEN Government Data Act makes all non-sensitive U.S. government data freely available and accessible to the public. Intel supports this bill and calls for its swift passage.

Second, the U.S. Government can help prepare an AI workforce. Workers need to have the right skills to create AI technologies and right now we have too few workers to do the job. Supporting universal STEM education is a start, but federal funding for basic scientific research that flows through agencies like the National Science Foundation to universities and helps to both train graduate level scientists and contribute to our scientific knowledge base on which industry can build are key. Current federal funding levels are not keeping pace with the rest of the industrialized world. I encourage lawmakers to consider the tremendous returns on investment to our economy that funding scientific research produces.

In addition to developing the right talent to develop AI solutions, governments will have to deal effectively with displaced workers. AI's emergence will displace some workers. But too little is known about the types of jobs and industries that will be affected most. Bills like H.R. 4829, the AI JOBS Act, that was recently introduced, help bridge that information gap by calling for the Department of Labor to study the issue and to work with Congress on recommendations. Intel supports this bill as well, and encourages Congress to consider it in Committee.

Finally, while AI holds many benefits, the government should play a role in helping society transition as AI emerges as a technology to mitigate undesirable impacts. Partnering with industry, government should set up guidelines to encourage ethical deployment of AI and to prevent it from being used in improper ways that would hurt the public. AI should not exist in "black boxes" – closed systems that receive an input, produce an output, and offer no explanation as to why or how⁸ – especially as we

⁸ <https://ai.intel.com/the-challenges-and-opportunities-of-explainable-ai/> (last accessed January 29, 2018).

begin to rely on AI to help us with things like criminal sentencing, determining insurance policies, or the job candidate referral processes.

Ensuring transparent decisions, predictability, resistance to manipulation, and accountability will be important as we strive for “AI for good”. Being able to explain why an algorithm made the decision it did will allow businesses and customers to make more informed choices about the AI tools that they use, and enable us to assign responsibility for outcomes to the proper human stakeholders. For these reasons, government should consider which AI uses require algorithm explainability to mitigate any potential for discrimination and harm to individuals.

As U.S. lawmakers consider whether and how – from legislative or regulatory standpoint – to deal with the emergence of AI, we encourage you to consider that technology like Artificial Intelligence which can enable innovation, economic growth and solve serious social issues, can only do so with a public policy environment that fosters its development. The United States has a long history of creating such an environment and to the extent it can continue to do so, it will reap the benefits of the technologies we are creating. Other countries see these potential benefits and are investing both in research and development and in creating favorable public policy environments. Intel recommends to all countries that they follow the proposals we include in our AI white paper, which I included as an addendum to my testimony. It is not a zero sum game and productive public policy environments in many countries can reinforce each other. The United States should specifically look at the proposals to make certain it does not fall behind.

Thank you again for the opportunity to testify today. We believe that AI is poised to create tremendous economic value while solving some of society’s most pressing challenges, but governments must mitigate unwanted impacts. Striking the right balance will require governments to embrace changes and help society prepare. This will be an iterative process. As you consider how existing legal and policy frameworks may need to be altered to accomplish this, I hope you’ll use Intel and other industry partners as a resource to offer technical expertise.

Thank you.



ADDENDUM

Artificial Intelligence The Public Policy Opportunity

Intel and Artificial Intelligence

Intel powers the cloud and billions of smart, connected computing devices. Due to the decreasing cost of computing enabled by Moore's Law⁹ and the increasing availability of connectivity, these connected devices are now generating millions of terabytes of data every day. Recent breakthroughs in computer and data science give us the ability to timely analyze and derive immense value from that data. As Intel distributes the computing capability of the data center across the entire global network, the impact of artificial intelligence is significantly increasing. Artificial intelligence is creating an opportunity to drive a new wave of economic progress while solving some of the world's most difficult problems. This is the artificial intelligence (AI) opportunity. To allow AI to realize its potential, governments need to create a public policy environment that fosters AI innovation, while also mitigating unintended societal consequences. This document presents Intel's AI public policy recommendations.

What is Artificial Intelligence?

While artificial intelligence is often equated with science fiction, it is not something looming on the horizon. It is already here, all around us, from the commonplace (talk-to-text, web searches, photo tagging, fraud detection) to the cutting edge (precision medicine, injury prediction, autonomous cars). Encompassing compute methods like advanced data analytics, computer vision, natural language

The **Collaborative Cancer Cloud**, a partnership between Intel, Oregon Health & Science University, Ontario Institute for Cancer Research and the Dana-Farber Cancer Institute, is enabling institutions to use distributed machine learning to speed up the discovery of new variants and biomarkers associated with cancer progression.

⁹ According to Moore's Law, the number of transistors on a chip roughly doubles every two years. As a result, the scale gets smaller and transistor count increases at a regular pace to provide improvements in integrated circuit functionality and performance while simultaneously decreasing costs.



processing, semantic graphs, and machine and deep learning, AI is transforming the way businesses operate and how people engage with the world.

While there isn't a commonly accepted definition for AI, Intel views it as a computerized system that performs tasks we normally associate with people. But in spite of the remarkable advances of computing power and sophisticated algorithms, there is still a long way to go before what is called General AI becomes a reality. General AI refers to the ability of a computerized system to portray human-like intelligence across a multitude of tasks.

In contrast, Narrow AI, which addresses a specific task or set of tasks, is commonplace and its use in many sectors of society will only increase. These are the technologies that help scientists gain novel insights into cancer diagnosis and aid in the design of new therapies, help physicians identify risks and predict the onset of diseases, allow interacting with our phones or vehicles' navigation systems through speech recognition, power household cleaning robots, help financial institutions fight against fraud, assist us with driving our cars and make manufacturing safer and more productive. Many other examples exist and the important point is that AI will transform many aspects of our lives.

Before considering the public policy impact of AI, it is crucial to introduce "machine learning," an important technique behind many of the recent advances of AI. Machine Learning is the development and application of algorithms to build and continually improve models from data. It allows computers to "learn" without being explicitly programmed. It is particularly useful when it is hard for humans to explain their innate ability to infer one thing from another; for example, how to distinguish a cat from a dog? Or how to understand speech? It is also useful when the amount of data is enormous in relation to human reasoning abilities (in examples such as ranking web pages through a search engine). By providing large datasets as input (for example, millions of videos) the computer starts to recognize patterns such as images of cats, without anyone instructing the computer on what a cat looks like. Another example is machine learning's ability to analyze many spam email messages, thereby "learning" to differentiate normal mail from spam.

Intel is a founding partner of **Hack Harassment**, a cooperative effort with the mission of reducing the prevalence and severity of online harassment. The initiative is evaluating AI technology and is working to develop an intelligent algorithm to detect and deter online harassment.

The Australian Government has funded an AI solution known as **Nadia**, an online virtual assistant who can speak, write, and chat online to interact with disabled persons. Nadia uses actress Cate Blanchett's voice and is programmed to improve its knowledge while answering thousands of questions asked in many different ways by people with diverse intellectual capabilities.



To take advantage of AI, all stakeholders must engage to understand the technology, debate how it will impact society and address concerns, as well as amplify its benefits and help society adjust.

Public Policy Considerations

The main drivers of public policy towards AI should be solving large societal problems and fostering economic progress. Accordingly, public policy must support industry efforts to bring AI benefits to the economy, to address citizens' concerns, and to identify needs for regulatory intervention.

As AI innovation is just beginning, it is crucial now to shape the public policy environment. Oversight by regulators will be essential for society to *trust* AI. Public policy should lower or remove any barriers standing between AI and its enormous potential to benefit our lives. Industry collaboration through groups like the Partnership on Artificial Intelligence¹⁰ are important to further study the issues and develop specific solutions.

Intel proposes the following AI public policy principles, and offers specific recommendations for government implementation:

- **Foster Innovation and Open Development** – To better understand the impact of AI and explore the broad diversity of AI implementations, public policy should encourage investment in AI R&D. Governments should support the controlled testing of AI systems to help industry, academia, and other stakeholders improve the technology.
- **Create New Human Employment Opportunities and Protect People's Welfare** – AI will change the way people work. Public policy in support of adding skills to the workforce and promoting employment across different sectors should enhance employment opportunities while also protecting people's welfare.
- **Liberate Data Responsibly** – AI is powered by access to data. Machine learning algorithms improve by analyzing more data over time; data access is imperative to achieve more enhanced AI model development and training. Removing barriers to the access of data will help machine learning and deep learning reach their full potential.

¹⁰ <https://www.partnershiponai.org/#>



- **Rethink Privacy** – Privacy approaches like The Fair Information Practice Principles and Privacy by Design have withstood the test of time and the evolution of new technology. But with innovation, we have had to “rethink” how we apply these models to new technology.
- **Require Accountability for Ethical Design and Implementation** – The social implications of computing have grown and will continue to expand as more people have access to implementations of AI. Public policy should work to identify and mitigate discrimination caused by the use of AI and encourage designing in protections against these harms.

Foster Innovation and Open Development

The potential of AI is enormous.¹¹ AI can enhance human capabilities, automate tedious or dangerous tasks keeping humans in safer conditions, unleash scientific discovery, and alleviate challenging societal problems. Doctors will be able to diagnose conditions earlier and more accurately, leading to quicker treatments and lives saved.¹² Automated vehicles will result in safer driving, and more efficiency and productivity. Farmers will increase crop yield based on real-time insights from weather and soil data, producing higher yields and more stable food supply even in unpredictable climates.

Realizing the potential of AI requires advances in core AI technologies. Governments must play a significant role in promoting those advances. Government investment in AI, public-private collaborations, and measures to incentivize adoption by society are public policy actions that will enable AI to develop and mature.

Equally, governments should gain expertise in AI in order to make effective public policy, to benefit from efficiency gains, and to champion AI adoption. Moreover, a new generation of AI specialists and data scientists should be on the radar of schools and universities when preparing new curricula.

Foster Innovation and Open Development – Recommendations

- **Fuel AI innovation:** Public policy should promote investment, make available funds for R&D, and address barriers to AI development and adoption.

¹¹ The global AI and robotics market is estimated to grow to \$153 billion by 2020 (Robot revolution – Global robot and AI primer, Bank of America Merrill Lynch, Dec 2015.)

¹² The market for AI system in healthcare is estimated to grow from \$633 million in 2014 to \$6 billion in 2021 (From \$600 M to \$6 billion, AI systems poised for dramatic market expansion in healthcare, Frost & Sullivan, Jan 2016.)



- **Address global societal challenges:** AI-powered flagship initiatives should be funded to find solutions to the world's greatest challenges such as curing cancer, ensuring food security, controlling climate change, and achieving inclusive economic growth.
- **Allow for experimentation:** Governments should create the conditions necessary for the controlled testing and experimentation of AI in the real world, such as designating self-driving test sites in cities.
- **Prepare a workforce for AI:** Governments should create incentives for students to pursue courses of study that will allow them to create the next generation of AI.
- **Lead by example:** Governments should lead the way on demonstrating the applications of AI in its interactions with citizens and invest sufficiently in infrastructure to support and deliver AI-based services.
- **Partnering for AI:** Governments should partner with industry, academia, and other stakeholders for the promotion of AI and debate ways to maximize its benefits for the economy.

Create New Employment Opportunities and Protect People's Welfare

Productive work is a fundamental component of individual wellbeing and high functioning societies. In the same way that AI needs to be designed to function properly, so should society be prepared to leverage AI's benefits while mitigating its impact on the workforce. While AI has the potential to improve many aspects of our lives and to spur economic growth, AI and robotics will bring automation to broad categories of jobs (e.g. fully autonomous vehicles will reduce the need for trucking and taxi drivers). Concurrently, new tasks and jobs will be created requiring entirely different sets of skills. Governments need to understand how AI will impact employment and have a plan to encourage employment in ways that allow technology to assist humans in the pursuit of their work.

From more timely, more accurate medical diagnostics to intelligent, safer transportation, AI will affect all facets of the economy, including the public sector. The economic benefits of AI should be inclusive, accessible, and broadly shared by society. Public policies must be enabled to mitigate inequalities, protect citizens' welfare, and help with the transition to a more data-driven economy.



Create New Employment Opportunities and Protect People's Welfare – Recommendations

- **Encouraging Human Employment:** Governments should implement programs to mitigate AI's impact on jobs and devise policies that promote employment. These programs should particularly focus on the effectiveness of incentives in government funded infrastructure projects.
- **Retraining:** Governments should implement policies that support the up-skilling and the re-skilling of the workforce, particularly in job areas that are less likely to be automated, such as positions focused on person to person interaction and the need for "guided computation" where individuals direct and oversee the operation of the technology.

Liberate Data Responsibly

AI does not exist without data. Machine learning based algorithms are trained with existing data and those data relate to specific usage domains. For instance, if AI is to be used to fight cancer, then deidentified data from medical records, genomic information, state of the art treatments and many other domains should be made available. Of particular interest are solutions that allow for the federated access to data from distributed repositories held in different sites, while preserving privacy and security.

Governments are also solicitors, creators and repositories of data. As long as no personal or sensitive information is involved, many of these datasets should be made available for public use. If personal or sensitive information is a requirement to solve critical societal problems (like making breakthroughs in personalized medicine), governments should partner with the private sector to find solutions to use AI while still delivering privacy protections. One example of such protection is the use of AI algorithms that analyze data in several encrypted yet separate datasets, but never require sharing of the data outside the encrypted area. These mechanisms for "raw data", instead of aggregated inferences, are much more useful for training data analytics necessary for AI. As explained before, AI requires data to function and public sector data is a valuable source of information to develop AI solutions to societal challenges.



Liberate Data Responsibly – Recommendations

- **Keep data moving:** Governments should eliminate unwarranted data localization mandates and enable secure international data transfers through international agreements and legal tools.
- **Open public data:** While protecting privacy, governments should make useful datasets publicly available when appropriate and provide guidance to startups and small and medium businesses for its reuse.
- **Support the creation of reliable data sets to test algorithms:** Governments should explore non-regulatory methods to encourage the development of testing data sets.
- **Federate access to data¹³:** Governments should partner with industry to promote AI tools to access encrypted data for analysis, while not requiring transfer of the data.

Rethink Privacy

Where the data used for AI originates from identifiable individuals, appropriate protections should be implemented to ensure that data is deidentified, lawfully accessed, processed, and kept safe. Robust privacy regulatory frameworks for the protection of personal data and cybersecurity should also apply to AI implementations. Intel is a proponent of technology neutral comprehensive privacy laws based on the Organization for Economic Cooperation and Development's Fair Information Practice Principles (the FIPPs), which are the global common language of privacy.

Intel has long supported and implemented Privacy by Design.¹⁴ Intel has previously demonstrated through our Rethinking Privacy¹⁵ project, that the FIPPs can be implemented during Privacy by Design processes to better protect individuals.

Questions may arise regarding the enforceability of privacy protections when a machine uses data autonomously. In these circumstances, accounting for privacy principles when designing technology will help protect individuals.

“Security Safeguards” is one of the FIPPs and is particularly critical in protecting the trustworthiness of AI implementations. AI can be used to foster both privacy and security by predicting the spread of cybersecurity attacks and helping organizations protect their data and

¹³ Instead of centralizing data from several institutions, federated access to data allows each institution to keep control of their data while enabling joint data analytics across all institutions.

¹⁴ *Privacy by Design* refers to the philosophy and approach of embedding privacy into the design specifications of various technologies.

¹⁵ <http://blogs.intel.com/policy/files/2015/01/RethinkingPrivacy.pdf>



AI algorithms/models. A critical component of allowing AI to better protect privacy and security will be the use of cybersecurity data to better predict future attacks. As the compute power of the data center is distributed across the entire network, the potential for AI to stop cyberattacks before they do significant harm will be greatly increased. This is one of many reasons why governments should promote the use and sharing of data for cybersecurity purposes.

Rethink Privacy – Recommendations

- **Adopt Robust Privacy Laws:** Based on the OECD Fair Information Practice Principles.
- **Implement Privacy by Design:** Follow Intel's Rethinking Privacy approach to implement Privacy by Design into AI product and project development.
- **Keep data secure:** Policies should help enable cutting-edge AI technology with robust cyber and physical security to mitigate risks of attacks and promote trust from society.
- **It takes data for AI to protect data:** Governments should adopt policies to reduce barriers to the sharing of data for cybersecurity purposes.

Require Accountability for Ethical Design and Implementation

Trust in AI requires organizations to demonstrate to the public and government regulating bodies that the technology is designed, implemented, and operated responsibly.

The Information Accountability Foundation (IAF¹⁶) has spent considerable time articulating the essential elements of what is required to demonstrate the responsible handling of information. The IAF's five principles are:

1. Organization commitment to accountability and adoption of internal policies consistent with external criteria.
2. Mechanisms to put privacy policies into effect, including tools, training and education.
3. Systems for internal ongoing oversight and assurance reviews and external verification.
4. Transparency and mechanisms for individual participation.
5. Means for remediation and external enforcement.

With only small adjustments (amending the word "privacy" in the second principle to cover broader categories of automated decision making), this work can and should apply more broadly to AI. Organizations that develop and implement AI solutions will benefit from working through the principles as the resulting policies, processes, and resources put in place will demonstrate responsible behavior to both regulators and individuals who are impacted by AI solutions.

¹⁶ <http://informationaccountability.org/>



Applying the principles to AI requires new thinking. As an example, transparency may be more difficult for some AI approaches than with traditional data processing. Some algorithms use hundreds of millions of adjustable parameters to function and may be continually updated based upon real-time data. In some cases this makes it impossible to deconstruct how a particular result was produced by the algorithm to accurately trace back a cause. In other words, it may be impossible to understand how a result is achieved, consequently making AI less accountable to the user. However, there is ongoing research to derive rules from deep neural networks, and these algorithms are being used successfully, for example for sensory recognition (like image recognition and natural language speech interfaces) and fraud detection by financial institutions. There are also new approaches to AI such as memory-based reasoning that can better warrant outputs, including reference back to relevant prior episodes of similar persons, places, and things as raw evidence in the explanations to the user. More symbolic approaches to AI also claim transparency in backtracking of inferential logic chains.

Ensuring fairness of AI results depends upon how the algorithms were developed, and in the case of AI-based machine learning, also upon the data that was utilized for their training. Noting that AI algorithms have the potential to make less biased decisions than people, there is still a risk for unintended bias, and therefore unintended discrimination of individuals. This may happen, for example, when the data used to train the algorithm was not representative of the problem space in question. One example of this situation could be when the training datasets were not free from bias themselves. Means to mitigate bias include using algorithms and data models that account for bias, well-curated training sets, extensive verification and validation of AI systems, and alertness to possible ethical or fairness implications from AI-based decisions. Government and the private sector should continue to work together to study and develop solutions to regulate discrimination caused by AI implementations.

Require Accountability for Ethical Design and Implementation – Recommendations

- **Standing for “Accountable Artificial Intelligence”:** Governments, industry and academia should apply the Information Accountability Foundation’s principles to AI. Organizations implementing AI solutions should be able to demonstrate to regulators that they have the right processes, policies and resources in place to meet those principles.
- **Transparent decisions:** Governments should determine which AI implementations require algorithm explainability to mitigate discrimination and harm to individuals.



SUMMARY OF RECOMMENDATIONS

FOSTER INNOVATION AND OPEN DEVELOPMENT

- **Fuel AI innovation:** Public policy should promote investment, make available funds for R&D, and address barriers to AI development and adoption.
- **Address global societal challenges:** AI-powered flagship initiatives should be funded to find solutions to the world's greatest challenges such as curing cancer, ensuring food security, controlling climate change, and achieving inclusive economic growth.
- **Allow for experimentation:** Governments should create the conditions necessary for the controlled testing and experimentation of AI in the real world, such as designating self-driving test sites in cities.
- **Prepare a workforce for AI:** Governments should create incentives for students to pursue courses of study that will allow them to create the next generation of AI.
- **Lead by example:** Governments should lead the way on demonstrating the applications of AI in its interactions with citizens and invest sufficiently in infrastructure to support and deliver AI-based services.
- **Partnering for AI:** Governments should partner with industry, academia, and other stakeholders for the promotion of AI and debate ways to maximize its benefits for the economy.

CREATE NEW EMPLOYMENT OPPORTUNITIES AND PROTECT PEOPLE'S WELFARE

- **Encouraging Human Employment:** Governments should implement programs to mitigate AI's impact on jobs and devise policies that promote employment. These programs should particularly focus on the effectiveness of incentives in government funded infrastructure projects.
- **Retraining:** Governments should implement policies that support the up-skilling and the re-skilling of the workforce, particularly in job areas that are less likely to be automated, such as positions focused on person to person interaction and the need for "guided computation" where individuals direct and oversee the operation of the technology.



LIBERATE DATA RESPONSIBLY

- **Keep data moving:** Governments should eliminate unwarranted data localization mandates and enable secure international data transfers through international agreements and legal tools.
- **Open public data:** While protecting privacy, governments should make useful datasets publicly available when appropriate and provide guidance to startups and small and medium businesses for its reuse.
- **Support the creation of reliable data sets to test algorithms:** Governments should explore non-regulatory methods to encourage the development of testing data sets.
- **Federate access to data¹⁷:** Governments should partner with industry to promote AI tools to access encrypted data for analysis, while not requiring transfer of the data.

RETHINK PRIVACY

- **Adopt Robust Privacy Laws:** Based on the OECD Fair Information Practice Principles.
- **Implement Privacy by Design:** Follow Intel's Rethinking Privacy approach to implement Privacy by Design into AI product and project development.
- **Keep data secure:** Policies should help enable cutting-edge AI technology with robust cyber and physical security to mitigate risks of attacks and promote trust from society.
- **It takes data for AI to protect data:** Governments should adopt policies to reduce barriers to the sharing of data for cybersecurity purposes.

REQUIRE ACCOUNTABILITY FOR ETHICAL DESIGN AND IMPLEMENTATION

- **Standing for "Accountable Artificial Intelligence":** Governments, industry and academia should apply the Information Accountability Foundation's principles to AI. Organizations implementing AI solutions should be able to demonstrate to regulators that they have the right processes, policies and resources in place to meet those principles.
- **Transparent decisions:** Governments should determine which AI implementations require algorithm explainability to mitigate discrimination and harm to individuals.

¹⁷ Instead of centralizing data from several institutions, federated access to data allows each institution to keep control of their data while enabling joint data analytics across all institutions.