3 May 2024

BSA COMMENTS TO SELECT COMMITTEE ON ADOPTING ARTIFICIAL INTELLIGENCE

Submitted Electronically to the Select Committee on Adopting Artificial Intelligence (AI)

BSA | The Software Alliance (**BSA**)[1] welcomes the opportunity to submit comments to the Select Committee on Adopting Artificial Intelligence (AI) (**Committee**).

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, security solutions, and collaboration software. Our members are on the leading edge of providing AI-enabled products and services, and tools used by others in the development of AI systems and applications in Australia and globally. As a result, they have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

We welcome the Committee's efforts to report on the opportunities and impacts arising from the uptake of AI technologies. In the course of engaging with various governments on this topic,[2] we have worked with our members to develop resources and position papers that showcase how AI-enabled tools can spread the benefits of digital transformation, as well as the policies needed to build trust and confidence in AI's responsible development and deployment.

We invite the Committee to consider the following BSA resources, which we have annexed to this cover letter:

1. **Annex A: BSA Policy Solutions for Building Responsible AI.** BSA's Policy Solutions for Building Responsible AI is a comprehensive set of recommendations for policymakers worldwide to address AI policy issues and advance the adoption of responsible AI across the economy. It lays out the issues that policymakers must address to facilitate responsible AI innovation, and

[1] BSA's members include: Adobe, Alteryx, Altium, Amazon Web Services, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, Cloudflare, CNC/Mastercam, Cohere, Dassault, Databricks, DocuSign, Dropbox, Elastic, ESTECO SpA, Graphisoft, Hubspot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Nikon, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, Prokon, Rockwell, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Workday, Zendesk, and Zoom Video Communications, Inc.

[2] BSA has testified before the United States Congress and the European Parliament on the Framework and its approach to mitigating AI-related risks. See:

1) Testimony of Victoria Espinel, Public Hearing on AI & Bias, Special Committee on Artificial Intelligence in a Digital Age, European Parliament, November 2021, https://www.europarl.europa.eu/cmsdata/244265/AIDA_Verbatim_30_November_2021_EN.pdf;

2) Testimony of Aaron Cooper, Task Force on Artificial Intelligence: Beyond I, Robot: Ethics, Artificial Intelligence and the Digital Age, before the House Financial Services Committee, October 2021, https://www.congress.gov/117/meeting/house/114125/witnesses/HHRG-117-BA00-Wstate-CooperA-20211013.pdf.

proposes a variety of policy solutions to the identified issues. These include: a) encouraging global harmonisation; b) implementing strong corporate governance practices to mitigate AI risks; c) promoting innovation and creativity; d) protecting privacy; e) facilitating government use and procurement of AI; f) promoting transparency; g) advancing cybersecurity with AI; h) protecting national security; i) promoting multiple development models; j) supporting sound data innovation policies; k) investing in AI research and development; and l) building the workforce for an AI future.

2. **Annex B: Submission on Supporting Safe and Responsible AI in Australia.** In our submission to the Department of Industry, Science and Resources (**DISR**), we provided recommendations responding to DISR's Discussion Paper on Safe and Responsible AI in Australia.[3] Key recommendations include: a) defining and distinguishing AI developers and AI deployers, and tailoring their respective obligations to their different roles in the AI ecosystem; b) focusing regulatory efforts on high-risk AI use cases; and c) incorporating the use of impact assessments for high-risk AI use cases.

3. **Annex C: Everyday AI for Businesses & Everyday AI for Consumers.** In Everyday AI for Businesses, we provide examples of business-to-business uses of AI, focusing on how companies use AI in low-risk ways to create significant benefits to both businesses and the customers they serve, such as improving logistics and document management. In Everyday AI for Consumers, we highlight how consumers rely on a wide array of services powered by AI in their day-to-day lives, such as auto-complete functions and virtual backgrounds on video calls.

4. **Annex D: AI and Copyright Policy**. This document sets out BSA's positions on three issues pertaining to the application of copyright laws to AI: a) responsible AI training and protecting copyright holders; b) remedies if AI-generated works infringe; and c) copyright protection for creators using AI.

We hope that our resources will assist the Committee in its inquiry. For more AI-related resources, we encourage the Committee to visit **BSA's AI Resource Center**, which serves as a single point of information for our policy submissions and material on AI issues.

We look forward to serving the Committee as a resource as you continue to engage in discussions on this issue. Please do not hesitate to contact me if you have any questions regarding this submission or if I can be of further assistance

Sincerely,

Tham Shen Hong

Tham Shen Hong
Senior Manager, Policy – APAC

[3] Safe and Responsible AI in Australia Discussion Paper, June 2023, https://storage.googleapis.com/converlens-au-industry/industry/p/prj2452c8e24d7a400c72429/public_assets/Safe-and-responsible-AI-in-Australia-discussion-paper.pdf.

Annex A: BSA Policy Solutions for Building Responsible AI

The Software Alliance

BSA

BSA Policy Solutions for Building Responsible AI

Artificial intelligence (AI)-enabled software is helping businesses in every sector of the economy leverage the value of data to drive digital transformation. From manufacturers that use AI to design more innovative products to small businesses that rely on automated translation capabilities to grow their global customer base, AI is creating new opportunities to solve complex challenges. BSA members are at the forefront of the responsible development of AI, providing trusted software solutions that enable enterprises and their customers to harness the power of AI to improve their product offerings and enhance their competitiveness in critical areas such as health care, defense and infrastructure, and education. Rapid advances in AI are transforming expectations about how the technology may reshape the world. However, unlocking the full potential of AI will require a dynamic and flexible policy framework that spurs responsible AI innovation and use through enhanced accountability and transparency.

BSA SUPPORTS:

Global Harmonization

Reducing Risk Through AI Governance

Promoting Innovation and Creativity

Protecting Privacy

Facilitating Procurement and Government Use of AI

Promoting Transparency

Enabling Cybersecurity

Ensuring National Security

Promoting Multiple Development Models

Supporting Sound Data Innovation Policies

Investing in Research and Development

Investing in Workforce Development

Global Harmonization

Policymakers around the world are developing regulatory approaches to AI. The global nature of today's technology ecosystem demands coordinated policy responses to foster innovation.

BSA SUPPORTS

» **Pursuing interoperability.** Countries should work together to promote multistakeholder dialogue and develop a shared vision for a risk-based policy approach for addressing common AI challenges and advancing norms around responsible AI governance (e.g., risk-based approach to regulation, balanced responsibilities along the AI value chain). Global partners should also agree on common AI terminology and taxonomy, including building on ongoing work in the EU-US Trade and Technology Council.

Reducing Risk Through AI Governance

BSA members are advancing trust and ethics in AI and investing in research and development to address some of society's most pressing challenges. Organizations should ensure that society can realize the benefits of AI by proactively addressing its risks. A range of corporate governance safeguards can promote accountability by helping to identify and mitigate such risks and appropriately delineating roles and responsibilities along the AI value chain.

BSA SUPPORTS

» **Implementing risk management programs.** BSA supports implementing risk management programs to enable organizations to identify the personnel, policies, and processes necessary to manage AI risks. Elements of a risk management program may include clearly assigning roles and responsibilities, establishing formal policies, using evaluation mechanisms, ensuring executive oversight, performing impact assessments for high-risk AI, and having internal independent review mechanisms, such as interdepartmental governance or ethics committees, to evaluate and address AI issues that pose high risks. Organizations can incorporate these practices as part of a broader corporate risk management program or as a separate AI program.

» **Requiring impact assessments for high-risk uses of AI.** An impact assessment is an accountability mechanism that promotes trust by demonstrating a system has been designed and deployed in a manner that accounts for potential risks it may pose to the public. By establishing a process for personnel to document key design and deployment choices and their underlying rationale, impact assessments enable organizations to identify and mitigate risks that can emerge throughout a system's life cycle. BSA supports requirements for organizations that develop or deploy high-risk AI to conduct impact assessments and publicly affirm that they have complied with this practice. An AI system may be high-risk if it makes consequential decisions that determine an individual's eligibility for and result in the provision or denial of housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance.

» **Distinguishing between different actors in the AI ecosystem.** Obligations should be placed on organizations based on their role in the AI ecosystem so that they can appropriately address the risks that fall within their responsibilities. For example, an AI developer, an AI deployer, and other parties within the value chain will have different information about how the AI system was developed or operates, and the law should recognize these distinctions.

» **Testing high-risk AI systems.** BSA encourages measures that incentivize safety and security. Robust testing and evaluation of high-risk AI systems for safety, security, accuracy, and fairness is critical and is prioritized in the NIST AI Risk Management Framework, which BSA supports. Existing technical standards for AI testing are nascent and should be developed consistent with longstanding voluntary, market-driven, and consensus-based approaches to standards development.

- **Ensuring appropriate policies and information sharing for foundation models.** Any public policies regulating foundation models should be commensurate with the models' risks and capabilities. Foundation model developers should provide information about model capabilities, limitations, testing, and security along the AI value chain based on the level of risk involved.

Promoting Innovation and Creativity

AI is advancing innovation and creativity in every sector of the economy. As this technology continues to evolve, it is important to consider the role of copyright law in both encouraging innovation and protecting the rights of creators. Copyright law is sufficiently flexible to adapt to this transformational technology, but BSA encourages policymakers to consider whether additional protections are warranted to prevent the spread of unauthorized, AI-generated replicas of an artist's name, image, likeness, or voice.

BSA SUPPORTS

- **Recognizing the copyrightability of works created with the assistance of AI.** AI can bolster creativity, just as other software applications have long been an important tool of artists and storytellers (e.g., photo enhancements for visual artists, visual effects in media and entertainment, and arranging music for sound recordings). Copyright plays a key role in businesses' ability to protect creative material, including software code. The use of AI should not prevent a work developed in conjunction with human creativity from being eligible for copyright protection. If copyright protection is not available simply because AI was used in the creative process, it will limit the responsible use of AI and the purpose of copyright laws. As a result, the portions of the work that are influenced by human creativity should be protected by copyright laws. Lack of copyright protection may also cause innovators to seek out jurisdictions with laws and policies that are more protective of intellectual property.
- **Adopting voluntary methods for rights holders to opt out of AI training.** Access to sufficient data for training is critical to develop AI offerings that are as accurate and insightful as possible and optimize the benefits they can provide for organizations and society. In general, AI training involves computational analysis of data to identify probabilities, correlations, and trends, a process that does not typically use any of the data for its expressive content and, therefore, does not infringe any copyright in the underlying data. However, to support artists and rightsholders, BSA encourages industry to lead the development of automated tools to indicate that a rights-owner does not want a website used for training purposes, similar to the current "do not crawl" tools that apply to search engines.
- **Recognizing the sufficiency of existing copyright law to remedy infringement.** It is important to recognize that existing copyright law is sufficient to address when a work created with the assistance of AI infringes copyrighted material.
- **Enacting legislation to help protect content creators.** BSA supports developing legislation to afford public figures, musicians, singers, actors, and other creators with a right to prevent the unauthorized dissemination or use of their name, image, likeness, or voice and the unauthorized impersonation of creators in a manner consistent with First Amendment protections.

BSA supports requirements for organizations that develop or deploy high-risk AI to conduct impact assessments and publicly affirm that they have complied with this practice.

Protecting Privacy

The data-intensive nature of AI underscores the importance of meaningful consumer privacy protections. Consumers deserve to know how their personal data is used and protected, and consumer expectations should be backstopped by strong legal obligations on companies that collect or process personal information.

BSA SUPPORTS

- **Adopting comprehensive consumer privacy laws.** BSA supports comprehensive consumer privacy laws that establish strong consumer rights in their personal data, impose clear requirements on companies that handle that consumer data, provide robust security, promote the use of data for legitimate business purposes, and are backed by robust government enforcement.
- **Providing targeted opt-outs for profiling.** Consumers should have the right to opt out of consequential, automated decisions that are made solely by AI without human interaction and that have a legal or similarly significant effect on individuals.
- **Developing privacy-enhancing technologies.** BSA encourages the development of privacy-enhancing technologies to strengthen AI safeguards. BSA recognizes that automation also plays a vital role in enabling cybersecurity tools that support data privacy. For example, AI is leveraged to help protect businesses against data breaches; protect data, devices, and networks; prevent unauthorized access to data; and improve an organization's recovery time, even after a data breach.

Facilitating Procurement and Government Use of AI

Governments leverage AI to fulfill important functions. In doing so, governments should ensure that they have access to the most advanced IT solutions and that they have processes to govern the responsible development and use of AI.

BSA SUPPORTS

- **Implementing the NIST AI Risk Management Framework (RMF).** The NIST AI RMF is a flexible framework that can help organizations govern, map, measure, and manage AI risks. Government agencies should follow the practices set forth by the NIST AI RMF, including for procurement purposes.
- **Pursuing multi-cloud procurement.** Government agencies should work with multiple cloud providers to leverage the breadth of innovation occurring across the cloud industry. Agencies should not put all their data in one cloud infrastructure, but rather leverage multiple cloud service providers' compute, AI, and other technologies. BSA supports agencies using multi-cloud in cloud purchasing.
- **Deploying AI to meet today's challenges with today's solutions.** Governments must invest in AI-driven cybersecurity solutions to bolster their defenses and keep pace with malicious actors who are already using AI to improve their exploits.
- **Enabling use of commercial sector AI applications.** AI applications are and will continue to be built into commercial software, including that procured by governments. Governments should continue to prioritize adopting commercial software and embrace trustworthy AI solutions contained within it to enhance citizen services and improve operations. This requires governments to ensure AI policies do not inadvertently prevent the government from adopting low-risk commercial AI applications.

Promoting Transparency

There has been tremendous innovation in AI, but it can also exacerbate risks of misinformation. Transparency about AI-generated content is key to ensuring responsible AI.

BSA SUPPORTS

» **Encouraging the use of watermarks or other disclosure methods for AI-generated content.** These disclosures can help consumers tell whether content is human- or AI-generated. This can be helpful in preventing misinformation. Encouraging the use of watermarks or other disclosure methods for AI-generated content can help address this concern.

» **Promoting the Coalition for Content Provenance and Authenticity standard.** BSA supports the Content Authenticity Initiative's (CAI) efforts to promote the open Coalition for Content Provenance and Authenticity standard for content authenticity and provenance. This standard will help consumers decide what content is trustworthy and promote transparency around the use of AI. In conjunction with watermarking, the CAI approach provides secure, indelible provenance.

» **Disclosing when consumers are interacting with AI.** Consumers should know when they are interacting with AI depending on the circumstances and context of use. For example, chatbots should disclose that consumers are interacting with AI instead of a human. AI vendors should be prepared to provide some measure of explainability around models and outcomes.

Enabling Cybersecurity

Strong cybersecurity risk management is critical to combatting security threats. Although malicious actors can exploit AI to create security risks, AI can also be used to dramatically enhance cybersecurity.

BSA SUPPORTS

» **Using AI to improve secure software development.** Software producers should leverage AI to improve the secure software development process, including by identifying and remediating vulnerabilities.

» **Harnessing AI to improve cybersecurity risk management.** Policymakers should ensure that cyber defenders can flexibly use AI to provide an accurate understanding of organizations' attack surfaces and improve threat detection and security outcomes.

Ensuring National Security

AI could have implications for national security. BSA recognizes the need for targeted actions to protect against AI-related national security threats. Failure to proactively develop robust AI policies at the national level could create substantial gaps in national security.

BSA SUPPORTS

» **Using narrowly tailored measures to address national security risks.** BSA supports narrowly tailored efforts to protect national security that do not unnecessarily interfere with companies' ability to conduct routine business transactions.

» **Leveraging AI to improve critical infrastructure.** BSA recognizes that AI can contribute significantly to developing and improving critical infrastructure, such as transportation. BSA supports policies that facilitate the use of AI to enhance critical infrastructure. Efforts to mitigate risks to critical infrastructure should focus on instances where there is a risk that an AI system could override human control and endanger the health and safety of individuals.

Rules that unnecessarily limit cross-border data transfers or require data localization invariably limit the insights and other benefits that AI systems can provide.

Promoting Multiple Development Models

Open source is a critical component of the AI ecosystem. It expands the AI marketplace, enhances the diversity of product offerings, promotes transparency, and enables vulnerabilities to be identified and remediated.

BSA SUPPORTS

» **Continuing the development of open source AI.** AI policies should recognize the key role that open source plays in AI development. BSA encourages rules that support both open source and proprietary systems.

Supporting Sound Data Innovation Policies

The exponential increase in data, combined with increases in remote computing power and development of more sophisticated algorithms, has fueled progress in machine learning and AI.

Capitalizing on these developments to facilitate continued advances in AI requires sound data innovation policies.

BSA SUPPORTS

» **Facilitating global data flows.** Data transfers are integral to every stage of the AI life cycle, from developing predictive models to integrating and deploying AI systems. The data used in AI systems often originates from many geographically dispersed sources, making it imperative that data can move freely across borders. Rules that unnecessarily limit cross-border data transfers or require data localization invariably limit the insights and other benefits that AI systems can provide. In addition, countries should not require algorithmic disclosure as a condition for doing business.

» **Continuing efforts to make public government data sets open and available in machine-readable digital formats.** Government-generated data is an important asset that can serve as a powerful engine for creating new jobs, promoting economic growth, and enabling innovation in AI-related technologies. Governments collect and generate vast quantities of data that offer unique insights into virtually every facet of the modern world. To enhance AI innovation, governments should continue to prioritize the release of high-value, non-sensitive government data.

Investing in Research and Development

Government research and development (R&D) spurs technological innovation that can drive long-term economic growth. Strategic investment in education, research, and technological development will be integral to developing AI technologies.

BSA SUPPORTS

» **Increasing investment in R&D.** Increased funding for R&D is essential to sparking innovation, growing high-paying jobs, and ensuring economic competitiveness.

» **Encouraging R&D cooperation.** Countries should work together to identify and support R&D challenges across borders.

Investing in Workforce Development

AI is helping to generate new jobs across industry sectors and augmenting the current workforce. AI may also impact existing jobs, and BSA supports job training and retraining programs to minimize negative impact on workforces. Countries must not only ensure that they have the STEM talent needed to develop AI innovations, but also prepare the broader workforce for a future in which virtually every job will involve an increased interaction with AI and other technologies and the need for digital skills.

BSA SUPPORTS

» **Improving access and support for STEM education.** Broadening educational opportunities, improving training programs, and ensuring the development of a diverse workforce is needed to help meet the demand for skilled STEM workers.

» **Expanding workforce training and alternative pathways.** Industry and government should invest in programs to support creating alternative pathways to the full range of AI careers; this includes those that enable workers to develop high-demand technology skills without the need for a bachelor's or graduate degree. Programs like apprenticeships, partnerships with community colleges, digital skills training and certifications, boot camps, and public service opportunities are all important gateways to helping new and mid-career workers develop in-demand skills.

ADDITIONAL RESOURCES

Artificial Intelligence & Copyright Policy

The Software Alliance

BSA

Artificial Intelligence & Copyright Policy

Advancing Technology and Creativity in the 21st Century Economy

Artificial Intelligence (AI) is advancing innovation and creativity in every sector of our economy. For example, AI provides creators with new tools to enhance their craft—in special effects in film, in sound mixing, in architectural planning, and in vehicular styling and design. As this technology continues to evolve, it is important to consider the role of copyright law in encouraging innovation and protecting the rights of creators. US copyright law is sufficiently flexible to adapt to this transformational technology, but we encourage further work on additional protections for artists to prevent the spread of unauthorized, AI-generated replicas of their name, image, likeness, or voice.

COPYRIGHT LAW AND AI INNOVATION

Responsible AI Training and Protecting Artists and Copyright Holders

Remedies if AI-Generated Works Infringe

Copyright Protection for Creators Using AI

Responsible AI Training and Protecting Artists and Copyright Holders

Training AI systems involves the computational analysis of large volumes of data. An AI system turns bits of data into tokens and maps how a token correlates with others. Computational analysis allows the AI system to predict what will come next.

Copyright protection applies broadly to almost any creative expression. Some of the data used to train an AI system may be part of a copyrighted work. But the training data is normally not used for its expressive content. Rather, the data is disassembled into smaller machine-readable units—or "tokens"—and then put through a computational analysis that involves mathematical calculations of probabilities, correlations, trends, and other patterns across millions or billions of tokens in a training data set.

An AI developer training a large language model, for instance, may use publicly available textual material (ranging from public, but copyright-protected essays to anonymous commentary on a website) to create a training data set. The use of the data is only to extract

www.bsa.org

DOWNLOAD NOW

Impact Assessments: A Key Part of AI Accountability

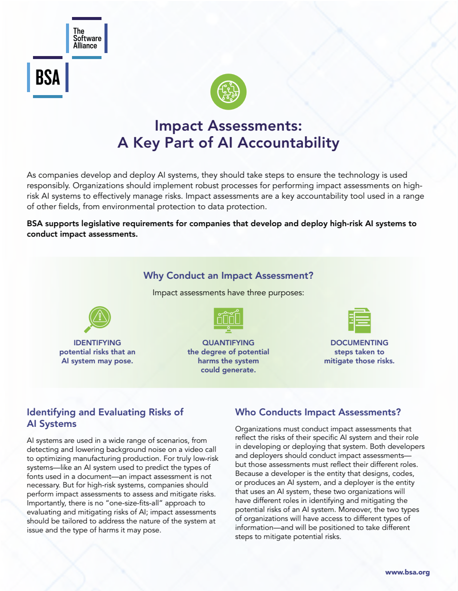

The Software Alliance

BSA

Impact Assessments:
A Key Part of AI Accountability

As companies develop and deploy AI systems, they should take steps to ensure the technology is used responsibly. Organizations should implement robust processes for performing impact assessments on high-risk AI systems to effectively manage risks. Impact assessments are a key accountability tool used in a range of other fields, from environmental protection to data protection.

BSA supports legislative requirements for companies that develop and deploy high-risk AI systems to conduct impact assessments.

Why Conduct an Impact Assessment?

Impact assessments have three purposes:

IDENTIFYING potential risks that an AI system may pose.

QUANTIFYING the degree of potential harms the system could generate.

DOCUMENTING steps taken to mitigate those risks.

Identifying and Evaluating Risks of AI Systems

AI systems are used in a wide range of scenarios, from detecting and lowering background noise on a video call to optimizing manufacturing production. For truly low-risk systems—like an AI system used to predict the types of fonts used in a document—an impact assessment is not necessary. But for high-risk systems, companies should perform impact assessments to assess and mitigate risks. Importantly, there is no "one-size-fits-all" approach to evaluating and mitigating risks of AI; impact assessments should be tailored to address the nature of the system at issue and the type of harms it may pose.

Who Conducts Impact Assessments?

Organizations must conduct impact assessments that reflect the risks of their specific AI system and their role in developing or deploying that system. Both developers and deployers should conduct impact assessments—but those assessments must reflect their different roles. Because a developer is the entity that designs, codes, or produces an AI system, and a deployer is the entity that uses an AI system, these two organizations will have different roles in identifying and mitigating the potential risks of an AI system. Moreover, the two types of organizations will have access to different types of information—and will be positioned to take different steps to mitigate potential risks.

www.bsa.org

DOWNLOAD NOW

The Software Alliance

BSA

www.bsa.org

Annex B: Submission on Supporting Safe and Responsible AI in Australia

26 July 2023

BSA COMMENTS ON SUPPORTING SAFE AND RESPONSIBLE ARTIFICIAL INTELLIGENCE IN AUSTRALIA

Submitted Electronically to the Department of Industry, Science and Resources

BSA | The Software Alliance (**BSA**)[1] welcomes the opportunity to submit comments to the Department of Industry, Science and Resources (**DISR**) on its Discussion Paper pertaining to Supporting Safe and Responsible AI in Australia (**Discussion Paper**).[2]

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, security solutions, and collaboration software. Our members are on the leading edge of providing AI-enabled products and services, and tools used by others in the development of AI systems and applications. As a result, they have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

We welcome DISR's efforts to explore how the Australian Government can support safe and responsible AI practices. As the Discussion Paper notes, AI presents significant opportunities and has delivered substantial benefits across both the economy and society, but also like any groundbreaking innovation, creates new risks and challenges.[3] BSA's views on how to best mitigate the risks of AI are informed by our recent experience working with member companies to develop the BSA Framework to Build Trust in AI (**BSA Framework**),[4] a risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle. Built on a vast body of research and informed by the experience of leading AI developers and deployers, the BSA Framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk mitigation best practices. BSA has testified before the United States Congress and the European Parliament on the Framework and its approach to mitigating AI-related risks.[5] Our extensive experience on these issues informs our response to the Discussion Paper.

[1] BSA's members include: Adobe, Alteryx, Altium, Amazon Web Services, Atlassian, Autodesk, Bentley Systems, Box, Cisco, Cloudflare, CNC/Mastercam, Dassault, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, IBM, Informatica, Juniper Networks, Kyndryl, MathWorks, Microsoft, Nikon, Okta, Oracle, Prokon, PTC, Rockwell, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

[2] Safe and Responsible AI in Australia Discussion Paper, June 2023, https://storage.googleapis.com/converlens-au-industry/industry/p/prj2452c8e24d7a400c72429/public_assets/Safe-and-responsible-AI-in-Australia-discussion-paper.pdf

[3] Discussion Paper (2023), p. 3 and p. 7.

[4] Confronting Bias: BSA's Framework to Build Trust in AI, June 2021, https://ai.bsa.org/wp-content/uploads/2021/06/2021bsaaibias.pdf, and enclosed.

[5] Testimony of Victoria Espinel, Public Hearing on AI & Bias, Special Committee on Artificial Intelligence in a Digital Age, European Parliament, November 2021, https://www.europarl.europa.eu/cmsdata/244265/AIDA_Verbatim_30_November_2021_EN.pdf; Testimony of Aaron Cooper, Task Force on Artificial Intelligence: Beyond I, Robot: Ethics, Artificial Intelligence and the Digital Age, before the House

Summary of BSA's Recommendations

BSA offers the following recommendations, which respond to the specific policy issues highlighted in the Discussion Paper.

Definitions

1. DISR should stay abreast of international efforts to define key terms, with a view to promoting harmonisation and interoperability.
2. AI policies should define, and consequently distinguish, AI developers and AI deployers. Both types of entities should have their respective obligations to ensure responsible AI innovation, and those obligations should be tailored to their different roles in the ecosystem.

Risk-based approaches

3. Obligations should fall on the entity that is best positioned to both identify and efficiently mitigate the risk of harm.
4. Regulatory efforts should be focused on high-risk AI use cases.

Strengthen risk management mechanisms

5. Impact assessments for high-risk uses of AI systems should play a significant role in AI risk management. BSA supports the development of AI policies that acknowledge the utility and incorporate the use of AI impact assessments.
6. DISR should draw on existing risk management frameworks, such as the US National Institute of Standards and Technology's (**NIST**) AI Risk Management Framework (**RMF**) and the BSA Framework and promote risk management tools to help Australian companies implement processes for managing risks of AI systems.

Target areas

7. To reduce the risk of AI bias, impact assessments should be conducted throughout the AI lifecycle.
8. Obligations aimed at enhancing transparency in AI systems should account for the different roles that AI developers and AI deployers play in the ecosystem.
9. Responsible AI innovation must be built on a solid foundation of data privacy. In the context of the proposed obligations on AI and ADM in the Privacy Act Review, BSA recommends: a) introducing a clear definition for "legal or similarly significant" that is interoperable with rights created in privacy laws internationally; and b) ensuring that these rights are exercised through the data controller and not the data processor.

Definitions

1. Stay abreast of international efforts to define key terms

The Discussion Paper defines the following terms: AI, Machine Learning, Generative AI models, Large Language Models (**LLMs**), Multimodal Foundation Model (**MfM**) and Automated Decision Making (**ADM**).

BSA is encouraged that, in defining these terms, DISR has referred to definitions found in other internationally recognised standards and frameworks. Specifically, the Discussion Paper's definitions of AI and Machine Learning are based on the International Standards Organization's (**ISO**) definitions.[6] In particular, the ISO definition of AI shares many similarities with the Organization for Economic Co-operation and Development's (**OECD**) definition of AI in their Recommendation of

Financial Services Committee, October 2021, https://www.congress.gov/117/meeting/house/114125/witnesses/HHRG-117-BA00-Wstate-CooperA-20211013.pdf.

[6] ISO/IEC 22989: 2022 – Artificial Intelligence Concepts and Terminology, July 2022, https://www.iso.org/obp/ui/#iso:std:iso-iec:22989:ed-1:v1:en.

Council on Artificial Intelligence (**Recommendation**). The NIST AI RMF adapts the OECD definition and defines AI as an "engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments" and that operates "with varying levels of autonomy".[7] We support this definition.

Given that AI systems are developed and deployed in an international context, definitions pertaining to AI should ideally be aligned across jurisdictions to ensure that all stakeholders have a common understanding of AI. As part of the work of the U.S.-EU Trade and Technology Council, the United States and the EU are agreeing on shared interpretations of key defined terms. For example, the EU-U.S. Terminology and Taxonomy for Artificial Intelligence defines machine learning as a branch of AI "and computer science which focuses on development of systems that are able to learn and adapt without following explicit instructions imitating the way that humans learn, gradually improving its accuracy, by using algorithms and statistical models to analyse and draw inferences from patterns in data."[8]

Prioritising international alignment in defining AI-related terms will: a) reduce discrepancies and conflicts between different legal frameworks, thus promoting compliance; b) serve as foundation for dialogue and cooperation between governments on AI-related risks; and c) support the international development of best practices and benchmarks for using AI systems safely, allowing AI systems to be deployed responsibly on a global scale.

Recommendation: DISR should stay abreast of international efforts to define key terms, with a view to promoting harmonisation and interoperability.

2. Define and distinguish "AI developers" and "AI deployers"

The Discussion Paper's definitions did not account for the different entities involved in an AI system's supply chain. Reflecting the inherently dynamic nature of AI systems, policies pertaining to AI must account for the array of stakeholders that may play a role in various aspects of a system's design, development, and deployment. In general, there are at least two key stakeholders with varying degrees of responsibility for managing the risks associated with an AI system throughout its lifecycle:

- **AI developers:** An AI developer is an entity that designs, codes, or produces an AI system.
- **AI deployers:** An AI deployer is an entity that uses an AI system.

It is crucial to define, and consequently distinguish, the AI developer and the AI deployer. As explained in the OECD's Recommendation, effective AI policies must necessarily account for "stakeholders according to their role and the context" in which AI is being deployed.[9] Effective management of risks among these different actors will depend on the nature of the AI system being developed. Distinguishing between AI developers and AI deployers ensures that specified obligations reflect an entity's role in the AI ecosystem. Tailoring obligations to an entity's role as an AI developer or AI deployer enables the company to fulfill the corresponding obligations and better protect consumers.

For example, an AI developer that designs an AI system is well-positioned to have access to information about the type of data that is used to train an AI system, the system's known limitations, and its intended use cases. However, the AI developer would *not* have insight into how the AI system is used after another company has purchased and deployed the AI system. Instead, the AI deployer – the entity using the AI system – is generally best positioned to provide details on how the system is

[7] NIST Artificial Intelligence Risk Management Framework 1.0, January 2023, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

[8] EU-U.S. Terminology and Taxonomy for Artificial Intelligence, May 2023, https://digital-strategy.ec.europa.eu/en/library/eu-us-terminology-and-taxonomy-artificial-intelligence.

[9] OECD Recommendation (2019). Per the Recommendation, the AI stakeholder community "encompasses all organizations and individuals involved in, or affected by, AI systems, directly or indirectly."

being used, the outputs from the AI system, the nature of any customer complaints, and other real-world factors affecting the system's performance. AI deployers are also best positioned to understand the risk profile that an AI system may present to individuals. Ensuring AI policies create obligations that reflect these different roles will enable all stakeholders to better understand how their organisations can identify and address harmful bias in AI systems.

Recommendation: AI policies should define, and consequently distinguish, AI developers and AI deployers. Both types of entities should have their respective obligations to ensure responsible AI innovation, and those obligations should be tailored to their different roles in the ecosystem.[10]

Risk-based approaches

3. Obligations should fall on appropriate entities

Obligations associated with AI should not apply to all actors in the AI ecosystem in the same way. BSA has emphasised the importance of distinguishing the AI developer from the AI deployer because, by implementing this distinction, it will minimise uncertainty about which actor will bear responsibility for complying with key legal requirements reflected in AI legislation or policies and better enable organisations to fulfill responsibilities according to their different roles.

Indeed, the regulatory requirements for the various entities responsible for developing and deploying AI should account for their unique roles and capabilities. Any obligation should fall on the entity that is best positioned to both identify and efficiently mitigate the risk of harm. We support AI governance approaches that promote accountability for both AI developers and deployers respectively, as each entity in the design, development and deployment of AI should have clear responsibilities.

For example, when using impact assessments to manage AI risks, both AI developers and deployers should document key aspects of AI systems, which are important reference points for understanding the operation of AI systems. However, the information to be documented will be different for developers that design an AI system than for deployers using an AI system:

- **Developers of high-risk AI systems** should document information including, as appropriate:
  - The intended purpose of the AI system;
  - Known limitations of the AI system;
  - Known, likely, and specific high risks that could occur and steps taken to mitigate those risks;
  - An overview of the data used to train the AI system; and
  - A summary of how the AI system was evaluated prior to sale.
- **Deployers of high-risk AI systems** should document information including, as appropriate:
  - The purpose for which the deployer intends to use the AI system;
  - Transparency measures, including notices to impacted individuals about the AI system's use;
  - A summary of how the AI system is evaluated, if applicable;
  - Known, likely, and specific high risks that could occur and steps taken to mitigate those risks; and
  - Post-deployment monitoring and user safeguards, if applicable.

Recommendation: Obligations should fall on the entity that is best positioned to both identify and efficiently mitigate the risk of harm.

[10] For further information, please see: AI Developers and Deployers: An Important Distinction, March 2023, https://www.bsa.org/files/policy-filings/03162023aidevdep.pdf

4. Focus on high-risk AI use cases

The AI ecosystem is broad, encompassing a diverse range of technologies and use cases and a wide array of stakeholders. Because the risks of AI are inherently use-case specific, any regulations should focus on specific applications of the technology that pose higher risks to the public but should be flexible enough to account for the unique considerations that may be implicated by specific use cases.

As a general principle, the scope of any regulatory obligations should be a function of the degree of risk and the potential scope and severity of harm. Many AI systems pose extremely low, or even no, risk to individuals or society, while creating potentially significant benefits. Imposing onerous regulations on such low-risk systems would hamper AI innovation with few corresponding benefits and therefore limit opportunities to use AI for positive impact. For example, AI is a critical component of cybersecurity risk mitigation,[11] which creates significant benefits to both companies and to consumers. Regulation should be mindful of the unintended consequences of regulations that could inadvertently limit the deployment of AI in detecting and responding to ever-evolving cybersecurity threats.

AI regulatory efforts should focus on addressing high-risk AI use cases. In this regard, AI systems may be high risk if they are used to make decisions to hire, promote, or terminate an individual's employment, or in other contexts, to determine eligibility for credit, healthcare, or housing.

Recommendation: Regulatory efforts should be focused on high-risk AI use cases.

Strengthen risk management mechanisms

5. Impact assessments should play a significant role in AI risk management

The Discussion Paper identified a range of AI-related risks, including concerns regarding algorithmic bias, misleading outputs, privacy, and transparency.[12] In view of these risks, accountability mechanisms, which enable organisations to demonstrate that they have developed or deployed AI responsibly, are crucial to maintain public confidence in AI systems. While the Discussion Paper highlighted various regulatory options,[13] BSA notes that the use of impact assessments was only briefly considered as part of a draft risk-based approach for AI.[14]

Impact assessments should play a significant role in DISR's approach to AI risk management. Impact assessments are an important accountability mechanism that exists today and can be applied to AI, as they can be used to help AI developers and deployers of AI systems for high-risk uses identify and mitigate risks throughout the lifecycle of an AI system. By allowing personnel across the organisation to examine the objectives, data preparation, design choices, and testing results, impact assessments help to drive internal changes to an organisation's risk management program. Implementing these changes enables organisations to better address existing concerns and adapt to new risks as they emerge. The fact that assessments are being performed for high-risk uses of AI systems also promotes trust for external stakeholders because they will know that an organisation is conducting a thorough examination of AI systems, and that the assessments are available to regulators upon request in the event of an investigation.

A recent report on AI accountability also concluded that impact assessments had several advantages over other accountability tools, noting that: 1) they are familiar to organisations already conducting impact assessments for privacy and data protection; 2) they are practical because they do not rely on technical standards, which are currently nascent; and 3) they are future-proof because they can adapt

[11] An organisation could face millions of indicators of compromise per day and security teams demand contextual awareness and visibility from across their entire environments. Cybersecurity providers that leverage AI can detect and respond to both known and unknown threats in real-time, with speed and scale to match.

[12] Discussion Paper (2023), p. 7-8.

[13] Discussion Paper (2023), p. 28 – 31.

[14] Discussion Paper (2023), p. 40.

as AI systems and AI governance evolve.[15]

Notably, the Discussion Paper indicates that final results of an impact assessment should be published and that external experts should conduct reviews for high-risk uses. This approach creates concerns, because one reason that impact assessments create a strong accountability tool is that they are conducted internally through confidential assessments. One of the goals of impact assessments is to drive internal changes, and organisations will likely conduct less thorough reviews to surface problems if the results ultimately will either be made public or be shared with third parties. In addition, impact assessments will likely contain confidential business information that organisations want to protect. Further, sharing sensitive information with third parties could implicate privacy and security concerns. Accordingly, we recommend that the process of conducting an impact assessment remain an internal exercise without the involvement of third parties or publication of the results.

For the reasons discussed above, BSA strongly supports the use of impact assessments to mitigate risks arising from high-risk uses of AI systems. The BSA Framework similarly recommends the use of impact assessments for high-risk uses of AI systems and is appropriate for organisations to implement to enhance AI accountability.

Recommendation: Impact assessments should play a significant role in AI risk management for high-risk use cases. BSA supports the development of AI policies that acknowledge the utility and incorporate the use of AI impact assessments for high-risk uses.

6. Promote and use risk management tools

BSA supports a risk-based approach for addressing potential AI risks. Risk management frameworks and tools developed in collaboration with stakeholders across the AI ecosystem can help drive consensus around processes for identifying, measuring, mitigating, and communicating about AI risks. DISR should draw on existing risk management frameworks, such as the NIST AI RMF and the BSA Framework, and promote risk management tools to help Australian companies implement processes for managing risks of AI systems.

In this regard, the BSA Framework identifies specific practices that AI developers and deployers can implement to ensure that AI is developed and used responsibly. For example, the BSA Framework highlights a range of issues that impact assessments should address, including identifying fairness metrics that will be used to assess bias in the AI system, ensuring that senior leadership has been briefed on potential high-risk AI systems, scrutinising training data for bias, and documenting how testing was performed.

BSA has also strongly supported NIST's work to develop its AI RMF, in line with our longstanding recognition that risk management is a key component of promoting trust in AI. The NIST AI RMF identifies specific practices that can be implemented to develop and deploy trustworthy AI. Both frameworks recommend:

- Consultation with a diverse group of stakeholders;
- Establishing processes to identify, assess, and mitigate risks;
- Defining individual roles and responsibilities to people throughout an organisation;
- Identifying metrics for evaluation;
- Evaluating fairness and bias;
- Maintaining post-deployment feedback mechanisms; and
- Establishing detailed plans for responding to incidents.

BSA recently prepared a crosswalk between the two frameworks, which illustrates the significant

[15] Impact Assessments: Supporting AI Accountability & Trust, January 2023, https://accesspartnership.com/impact-assessments-supporting-ai-accountability/

alignment between the two approaches.[16] These frameworks are appropriately voluntary, and the flexibility to utilise different components allows organisations to tailor the frameworks to their business needs.

Recommendation: DISR should draw on existing risk management frameworks, such as the NIST AI RMF and the BSA Framework, and promote risk management tools to help Australian companies implement processes for managing risks of AI systems.

Target Areas

7. Reducing the risk of bias in AI systems

As highlighted in the Discussion Paper, "[a]lgorithmic bias is often raised as one of the biggest risks or dangers of AI," involving "systematic or repeated decisions that privilege one group over another."[17]

The BSA Framework elaborates on how bias could occur at multiple stages in the development and deployment of an AI system.[18] For example, with respect to training data, there are risks of perpetuating historical biases reflected in the data or of sampling bias – where the data is misrepresentative of the population in which it will be used. The process of selecting the input variables (i.e., features) that the model will weigh as it is being trained is another critical decision point that can introduce bias. Even when sensitive demographic data is excluded, bias may be introduced if the system relies on features that are closely correlated to those traits, called proxies. Bias can also arise in various ways after a system has been deployed, including when the data used to train or evaluate an AI system differs materially from the population the system encounters when it is deployed.

One way to identify and address these risks in high-risk systems is conducting an impact assessment. For example, the BSA Framework recommends that in the data acquisition phase of the design of an AI model, companies should evaluate the representativeness of the data as part of conducting an impact assessment. To do so, a company can compare the demographic distribution of training data to the population in which the system will be deployed and assess whether there is sufficient representation of subpopulations that are likely to interact with the system. To mitigate issues that arise, companies can consider rebalancing the dataset with additional data or synthetic data, which involves oversampling data from underrepresented groups.

Similarly, in the data preparation and model definition phase, an impact assessment could include documenting a potential correlation between selected features and sensitive demographic attributes. For features that closely correlate to a sensitive class, companies can document the relevance to the target variable and the rationale for its inclusion in the model, consistent with laws prohibiting discriminatory actions. These practices, along with other measures, can help to prevent algorithmic discrimination while enabling innovation.

Recommendation: To reduce the risk of AI bias, impact assessments should be conducted to address issues throughout the AI lifecycle.

8. Enhancing transparency in AI systems

The Discussion Paper notes the importance of transparency in ensuring "appropriate accountability, risk mitigation and responsibility for liability is applied appropriately across AI vendors and buyers along the value chain."[19]

BSA supports enhancing transparency of AI systems. As policy approaches to enhance transparency

[16] Crosswalk Between BSA Framework to Build Trust in AI and NIST AI Risk Management Framework, April 2023, https://www.bsa.org/policy-filings/us-crosswalk-between-bsa-framework-to-build-trust-in-ai-and-nist-ai-risk-management-framework

[17] Discussion Paper (2023), p. 8.

[18] BSA Framework (2021), p. 4-7.

[19] Discussion Paper (2023), p. 8-9.

are developed, both AI developers and AI deployers play distinct but important roles in ensuring appropriate transparency for consumers. Notably, while AI developers are best positioned to understand the intended uses and limitations of an AI system, AI deployers are the entities that typically interact with consumers and decide when and how to deploy the AI system for a particular use. Their respective obligations should reflect this difference in roles – for example, an AI developer might be required to maintain documentation summarising how its AI system was evaluated prior to sale, whereas an AI deployer might be required to maintain documentation on post-deployment monitoring and user safeguards.

In regard to transparency, a developer might be obligated to make available to a deployer information reasonably necessary for the deployer to conduct an impact assessment, including information about the AI system's capabilities, limitations, and intended purpose. A deployer's transparency obligations might focus on the information provided to consumers about the use of an AI system.

Recommendation: Obligations aimed at enhancing transparency in AI systems should account for the different roles that AI developers and AI deployers play in the ecosystem.

9. Safeguarding privacy in AI systems

The data-intensive nature of AI makes privacy protections vital. The Discussion Paper observed that, while "[r]ich, large and quality data sets are a fundamental input to AI", "access to and application of these datasets have the potential for individuals' data to be used in ways that raise privacy concerns."[20] Indeed, responsible AI innovation must be built on a solid foundation of data governance that prioritises the security of user data and aligns with the public's expectation about how their personal information will be used. To engender trust in AI systems, consumers must have confidence in how their data is used and protected, and consumers' expectations should be backstopped by strong legal obligations on companies that collect and/or process sensitive data.

Relatedly, BSA provided comments to the Attorney-General's Department's Privacy Act Review Report 2022 (**AGD Report**), where we addressed the AGD Report's proposals to introduce several new obligations related to AI and automated decision making (**ADM**) that have "legal or similarly significant effect" on an individual's rights.[21] BSA supports introducing robust consumer rights and obligations on companies that handle data. However, we also highlighted the importance of: a) creating a comprehensive definition for the phrase "legal or similarly significant effects" that is interoperable with other privacy laws, so as to increase certainty for both individuals and companies when the related rights are available and allow them to better apply these protections across different jurisdictions; and b) individuals should be exercising these rights through a data controller, rather than a processor acting on behalf of a controller.[22]

Recommendation: Responsible AI innovation must be built on a solid foundation of data privacy. In the context of the proposed obligations on AI and ADM in the Privacy Act Review, BSA recommends: a) introducing a clear definition for "legal or similarly significant" that is interoperable with rights created in privacy laws internationally; and b) ensuring that these rights are exercised through the data controller and not the data processor.

Conclusion

We hope that our comments will assist DISR in its objective to ensure AI is used safely and responsibly. We look forward to serving as a resource as you continue to engage in policy discussions on this issue. Please do not hesitate to contact me if you have any questions regarding this

[20] Discussion Paper (2023), p. 8.

[21] BSA Comments on the Privacy Act Review Report 2022, April 2023, https://www.bsa.org/files/policy-filings/04102023bsaaupriv.pdf, p. 10-11.

[22] Data controllers are entities that decide how and why to collect information, which should be distinguished from data processors, which are entities that process collected personal information on behalf of other entities. The controller-processor distinction is similar to the AI developer-deployer distinction in that it distinguishes the consumer-facing entity (the AI deployer and the data controller) from the non-consumer facing entity (the AI developer and the data processor).

submission or if I can be of further assistance.

Sincerely,

Tham Shen Hong

Tham Shen Hong
Manager, Policy – APAC

Annex C: Everyday AI for Businesses and Consumers

BSA The Software Alliance

Everyday AI for Businesses

Companies across industries are adopting artificial intelligence (AI) systems to improve the products and services they offer to consumers. These business-to-business (B2B) uses of AI demonstrate the many ways that organizations are already using AI to serve individuals. In many cases, companies will use AI in low-risk ways that create significant benefits—not only to the businesses but to the customers they serve. Businesses use AI systems everyday for routine tasks including:

Answering customer questions. Businesses can offer customers 24/7 support through AI-powered chatbots that answer straightforward questions even when human customer service representatives are asleep. Bots can be programmed to address basic questions, instead of sending customers to FAQs.

Responding to frequent emails. Companies can set up AI systems to respond to common requests—like sending automatic responses to emails asking about the status of payment invoices.

Improving logistics and planning. AI systems can improve a company's ability to forecast supply-chain issues, optimize delivery routes, estimate arrival times for new shipments, and reduce their fuel and energy usage.

Identifying and managing common documents. Companies can use AI tools to read hand-written documents, identify a contract based on its format, and scan files for sensitive data that needs stricter care. AI systems can then create summaries of regular corporate reports, or generate new forms based on existing examples from frequently-used documents.

Improving cybersecurity. AI systems can sift through large volumes of information created by users of a company's IT network to forecast, detect, prevent and respond to threats. AI systems can also distill large amounts of data about security events into concrete actions to help companies secure their products and services.

Keeping shelves stocked. AI systems can forecast demand for products and redistribute them across a company's physical stores. AI systems can also detect early signs of supply chain issues and alert managers if inventory drops below certain levels.

Improving safety for corporate cars. AI systems can be trained to alert employees about anomalies in corporate cars that can indicate maintenance or safety issues.

Transcribing meetings—and identifying action items. Employees can spend more time collaborating if they use an AI system to transcribe their conversation, summarize decisions, and identify follow-up items, instead of requiring an employee to serve as note-taker.

AI ACROSS SECTORS

AI systems can be used in a range of industry-specific scenarios, many of which help companies improve existing products and services.

Transportation

AI systems can improve the efficiency of airlines, by helping to pinpoint causes of any slowdowns in the process of cleaning, refueling, and reloading an airplane. Detecting these delays early helps the airline mitigate their effect on passengers.

Manufacturing

AI design tools can optimize manufacturing processes, to reduce waste and improve products. This is true from early phases, where AI can help design and test new prototypes, to factory floors where AI systems can identify maintenance and quality-control issues.

Agriculture

Farmers use AI systems to analyze large volumes of weather and crop information, helping them monitor their crops, increase yields, and adjust to rain and drought conditions.

Construction

Companies use AI to streamline the process of designing and constructing new buildings. They can also create "digital twins" of real-life cities to understand environmental and other impacts of a proposed design.

A GLOBAL APPROACH

To use these everyday AI systems, organizations generally need to bring together data they collect from many different regions, including construction sites, factory floors, or farm fields that may be located worldwide. Analyzing these different data points helps AI systems provide more accurate and reliable information.

Helping Consumers Behind the Scenes

Companies frequently use AI systems to protect consumers using their products and services, often in ways that are designed to operate in the background. For example, a range of businesses use AI systems to improve fraud detection and to protect against cybersecurity threats.

Fraud detection. Banks and credit card companies can use AI-powered systems to better detect potential fraud in real time, including by setting rules for identifying suspicious wire transfer and credit card transactions. AI systems enable companies to monitor large amounts of customer transactions to find anomalies, including detecting unusual usage patterns for a consumer's online accounts. Banks can use that information to alert customers about potential security concerns, while reducing the amount of "false positives" that may block consumers from using their own credit cards.

Cybersecurity. AI helps organizations stay a step ahead of hackers by predicting potential attacks, mitigating attacks in real-time, managing access to resources, and encrypting sensitive data. For example, a company can use an AI system to identify malicious files and suspicious IP addresses that can be easily missed by humans due to the sheer volume. In some cases, AI systems can be used to forecast, detect, prevent, and respond to threats automatically. This helps companies secure the products and services that consumers use, and protect against potential threats.

The Software Alliance

BSA

Everyday AI for Consumers

Consumers already rely on a wide array of services powered by artificial intelligence (AI). Although these AI systems may not gain widespread attention, they illustrate how useful AI systems are in everyday life. Many of these uses present few risks to individuals while creating significant benefits, including helping organize our digital files and improve our communications with friends and family.

Do you want to pick up where you left off?

AI systems are frequently used to identify documents and other files that users recently worked on and may want to re-open. These systems can also help users locate and organize their files, such as suggesting that similar files be stored in similar locations.

Do you want to know more about that athlete?

AI systems are used to improve traditional analytics that power fantasy sports leagues, by combining inputs on sports players and teams with news articles and other sources. That creates detailed insights for sports fans, like hole-by-hole player predictions for golf tournaments.

Do you want to save time completing a form?

AI systems can auto-populate your shipping address when you order a package or create draft responses to forms that you've completed in the past.

Do you want to use a virtual background?

AI systems power the increasingly popular virtual backgrounds available on video calls. Providers use AI systems to identify the outline of an individual so the virtual background can appear in the appropriate place and follow a user as she moves across the virtual screen.

Did you forget an attachment?

For years, AI systems have been used by email providers to identify when a user may have forgotten to attach a document—and ask if something is missing.

Is it noisy during your video call?

If you join a video call from a crowded room, the video call provider may use an AI system to reduce the amount of background noise heard by others on the call—while making sure you still come through loud and clear.

How can you reach that savings goal?

AI systems can help you track your spending and budget goals, including analyzing your monthly spending habits and providing personalized recommendations for saving money.

Do you need an answer quick?

Users can interact with AI-powered chatbots to find the answers they need, instead of scrolling through an entire website or a lengthy FAQ. Chatbots can be programmed to point consumers to helpful information like a company's return policy or a list of store locations.

Annex D: AI and Copyright Policy

The
Software
Alliance

BSA

Artificial Intelligence & Copyright Policy

Advancing Technology and Creativity in the 21st Century Economy

Artificial Intelligence (AI) is advancing innovation and creativity in every sector of our economy. For example, AI provides creators with new tools to enhance their craft— in special effects in film, in sound mixing, in architectural planning, and in vehicular styling and design. As this technology continues to evolve, it is important to consider the role of copyright law in encouraging innovation and protecting the rights of creators. US copyright law is sufficiently flexible to adapt to this transformational technology, but we encourage further work on additional protections for artists to prevent the spread of unauthorized, AI-generated replicas of their name, image, likeness, or voice.

COPYRIGHT LAW AND AI INNOVATION

- Responsible AI Training and Protecting Artists and Copyright Holders
- Remedies if AI-Generated Works Infringe
- Copyright Protection for Creators Using AI

Responsible AI Training and Protecting Artists and Copyright Holders

Training AI systems involves the computational analysis of large volumes of data. An AI system turns bits of data into tokens and maps how a token correlates with others. Computational analysis allows the AI system to predict what will come next.

Copyright protection applies broadly to almost any creative expression. Some of the data used to train an AI system may be part of a copyrighted work. But the training data is normally not used for its expressive content. Rather, the data is disassembled into smaller machine-readable units—or "tokens"—and then put through a computational analysis that involves mathematical calculations of probabilities, correlations, trends, and other patterns across millions or billions of tokens in a training data set.

An AI developer training a large language model, for instance, may use publicly available textual material (ranging from public, but copyright-protected essays to anonymous commentary on a website) to create a training data set. The use of the data is only to extract unprotected information about the English language (i.e., the correlations, patterns, and relationships among "tokens" spanning the 26 letters of the English alphabet and 1 million English language words, as they appear in thousands of stock phrases, figures of speech, similes, metaphors, and common expressions).

The AI training will not infringe copyright. Nonetheless, at BSA we believe it is important to consider additional steps to protect the creativity of artists and rightsholders. One step is to encourage voluntary conversations around automated tools to indicate that the rights-owner does not want a website used for training purposes, similar to the current "do not crawl" tools that apply to search engines. BSA supports further discussions to arrive at effective, consensus-based technical mechanisms.

Remedies if AI-Generated Works Infringe

Copyright holders should have full and effective remedies when their rights are infringed. This principle applies equally to outputs generated using AI systems.

Copyright remedies have been effective to deter infringement and should remain so. In addition, Congress should consider whether there are steps it can take to better protect artists from the spread of unauthorized, AI-generated replicas of their name, image, likeness, or voice. Many states currently provide some protections through rights-of-publicity. To improve protection for artists and creators, BSA supports further development of federal legislation focused on AI-generated replicas.

Copyright Protection for Creators Using AI

Generative AI can bolster creativity, just as other software applications have long been an important tool of artists and storytellers (e.g., photo enhancements for visual artists, special effects in audio-visual works, and arranging music for sound recordings). When generative AI is used to enhance human creativity, the resulting work should be protected by copyright. If copyright protection is not available simply because AI was used in the creative process, it will limit the responsible use of AI and the purpose of our copyright laws.