ARTIFICIAL INTELLIGENCE: AUSTRALIA'S ETHICS FRAMEWORK

COMMENTS FROM BSA | THE SOFTWARE ALLIANCE
May 31, 2019

Introduction and Summary of Comments

BSA | The Software Alliance (**BSA**)[1] thanks the Department of Industry, Innovation and Science (**Department**) for this opportunity to comment on the discussion paper on *Artificial Intelligence: Australia's Ethics Framework* (**Framework**)[2] produced by CSIRO's[3] Data61.

BSA is the leading advocate for the global software industry before governments and in the international marketplace. Our members are at the forefront of software-enabled innovation that is fueling global economic growth, including cloud computing and artificial intelligence (**AI**) products and services. BSA's members include many of the world's leading suppliers of software and online services and have made significant investments in developing innovative AI solutions for use across a range of applications.

As leaders in AI development, BSA's members have unique insights into both the tremendous potential of AI and the governmental policies that can best support the responsible use of AI and ensure continued innovation. To that end, BSA has identified five pillars[4] that are essential to the development of responsible AI frameworks. These pillars, with which the Framework is broadly aligned, reflect the fact that both industry and government have important roles to play in promoting the benefits and mitigating the potential risks involved in the development, deployment, and use of AI:

1) **Building Confidence and Trust in AI Systems**: Highlighting industry efforts to ensure AI systems are developed in ways that maximize fairness, accuracy, data provenance, explainability, and responsibility.

2) **Sound Data Innovation Policy**: Promoting data policies that are conducive to the development of AI, including reliable legal mechanisms that facilitate cross-border data transfers, legal certainty for value-added services (e.g., text and data mining, machine learning), and enhanced access to non-sensitive government data.

[1] BSA's members include: Adobe, Akamai, Amazon Web Services, Apple, Autodesk, AVEVA, Baseplan Software, Bentley Systems, Box, Cadence, Cisco, CNC/Mastercam, DataStax, DocuSign, IBM, Informatica, Intel, MathWorks, Microsoft, Okta, Oracle, PTC, Salesforce, ServiceNow, Siemens PLM Software, Sitecore, Slack, Splunk, Symantec, Synopsys, Trend Micro, Trimble Solutions Corporation, Twilio, and Workday.

[2] Available at: https://consult.industry.gov.au/strategic-policy/artificial-intelligence-ethics-framework/

[3] The Commonwealth Scientific and Industrial Research Organisation.

[4] For more information about these pillars, please visit https://www.ai.bsa.org

3) **Cybersecurity and Privacy Protection**: Advocating for policies that strengthen enhanced security measures and respect informed consumer choices while ensuring the ability to deliver valuable tailored products and services.

4) **Research and Development**: Supporting investment in efforts that foster confidence and trust in AI systems, promote coordination and collaboration between industry and government, and help grow the AI workforce pipeline.

5) **Workforce Development**: Identifying opportunities for government and industry to collaborate on initiatives to prepare the workforce for the jobs of the future.

BSA appreciates the opportunity to contribute to the development of the Framework. In this regard, we offer the following comments and recommendations, which focus on important aspects for creating a secure and trusted AI ecosystem, to further enhance the Framework.

In summary, BSA recommends that the Framework incorporate:

- software security as a core principle (in addition to the existing eight core principles that have been articulated in the Framework);

- sound data innovation policy to promote cross-border data transfers as an explicit tool in the "toolkit for ethical AI"; and

- a clear articulation of the multiple stakeholders involved, and their respective roles and responsibilities, in the development, deployment, and use of AI.

Our comments and recommendations are set out in greater detail below.

Detailed Comments

A. Incorporating Software Security as a Core Principle

As presently drafted, the Framework does well to acknowledge important supporting pillars, such as data governance[5] and protection against data breaches,[6] as being crucial to ethical AI. We particularly commend the inclusion of privacy protection as a core principle, as this is key to increasing individuals' feeling of trust and empowerment when interacting with AI services and solutions.

The Framework can be further enhanced by including software security as an additional core principle. As AI and other digital technologies increasingly create a globally connected economy, ensuring that AI systems are designed to mitigate foreseeable security risks will be critical.[7] The

[5] See section 3 of the Framework.

[6] See section 3.2 of the Framework.

[7] In relation to this, BSA commends to the Department the various materials we have developed to promote cybersecurity, while protecting privacy and safety (available at https://bsacybersecurity.bsa.org). This includes *BSA's International Cybersecurity Framework* which provides a recommended model for a comprehensive national cybersecurity policy, and is intended to serve as a tool both for policymakers considering foundational cybersecurity legislation and for those examining gaps and shortfalls in existing policies.Copy available at: http://bsacybersecurity.bsa.org/report-item/bsa-international-cybersecurity-policy-framework/

Framework would therefore benefit from including considerations related to securing software throughout its lifecycle. Software-enabled capabilities have expanded from traditional computer programs and industrial control systems into AI and emerging technology. These include widely deployed sensors, smart appliances, connected vehicles, and robotic systems. It is therefore imperative that software developers, including those developing AI solutions and applications, ensure that software is built and maintained securely throughout its lifecycle. In this regard, BSA has published a *Framework for Secure Software*[8] that serves as a comprehensive benchmark for software security considerations.

BSA recommends that software security considerations be added to the core principles of the Framework, to emphasize the importance of secure systems and software as part of ethical AI.

B. Incorporating Sound Data Innovation Policy in the Toolkit for Ethical AI

The Framework rightly identifies that data "is a key component of AI."[9] Indeed, the exponential increase in data, combined with increases in remote computing power and development of more sophisticated algorithms, has fueled advances in machine learning and AI. AI systems are "trained" by ingesting enormous volumes of data. The benefits of AI are therefore dependent on the quantity and quality of data that is available for training. As a result, government policies affecting the ability to access and share data have a significant influence on AI development.

It is commendable that the Framework also addresses the need to balance consumer protection with pro-innovation policies.[10] Frameworks such as "data-sharing legislation"[11] and the "Consumer Data Right"[12] already exist in Australia and provide good fundamentals to ensure the domestic availability and protection of data.

However, the Framework would further benefit by explicitly recognizing that sound data innovation policies are a critical component of a "toolkit for ethical AI." In addition to the data policies noted above, the Framework should also incorporate an analysis that recognizes the uniquely important role that cross-border data transfers play in the development and use of AI, and that promotes the free flow of data across borders. The free flow of data is integral to every stage of the AI life cycle, from the development of predictive models to the deployment and use of AI systems. Data used in AI systems often originates from many geographically dispersed sources. Furthermore, as recognized in the Framework, many AI solutions used in Australia are developed internationally[13] and offered over the cloud. Therefore, it is imperative that unwarranted data localization mandates be avoided, and that data be allowed to move freely across borders in an interoperable and secure way. Rules that limit cross-border data transfers invariably limit the insights and other benefits that AI systems can provide.

In this regard, BSA commends Australia's participation in the APEC Cross Border Privacy Rules system (**CBPR**). We look forward to Australia's implementation of the CBPR in 2019. In addition,

[8] See more at https://www.bsa.org/reports/bsa-framework-for-secure-software.

[9] See section 2.1.3 of the Framework.

[10] See section 3.1 of the Framework.

[11] Mentioned in section 2.1.3 of the Framework.

[12] Mentioned in section 3.3 of the Framework.

[13] See, for example, the paragraph on "International coordination is crucial" on page 8 and section 7.1.2 on page 59 of the Framework.

Australia can look toward other bilateral and multilateral data transfer agreements and commitments, as well as encourage the inclusion of data transfer provisions in trade agreements.

BSA recommends that the Framework's "toolkit for ethical AI" include a section on "data innovation policy" to promote cross-border data transfers and to eliminate unwarranted data localization.

C. Roles and Responsibilities of Different Stakeholders in Mitigating Risks in the Development, Deployment, and Use of AI

BSA welcomes the Framework's emphasis on the need to take a risk-based approach when putting ethical principles into practice, and the need for impact and risk assessments.[14] However, the Framework's risk-based approach would benefit from a more formal recognition and discussion of the distinct stakeholders in the AI value chain, and the importance of their respective roles and responsibilities in promoting ethical AI (e.g., AI solution providers, entities that deploy and use AI, and end users). Much like software security, "ethical AI" requires a lifecycle approach to risk management, which includes anticipating and addressing risks that can arise when systems are designed, after they have been deployed, and when they are being decommissioned. There is no one-size-fits-all approach for managing AI lifecycle risks. Indeed, the best risk-management practices must be tailored to account for an AI system's development model and deployment context. Allocating the appropriate roles and responsibilities for managing AI risks must likewise account for these considerations.

The Framework identifies a "toolkit" of methods that can be used for "implementing ethical AI". For instance, the Framework suggests the use of impact and risk assessments and the adoption of monitoring and recourse mechanisms as potential tools for ensuring that an AI system is consistent with the core principles of the Framework (see also our recommendation in section D2 below on impact and risk assessments). For the avoidance of doubt, it would be helpful for the Framework to clarify that the toolkit is relevant not only to developers of AI but also to entities that deploy and use AI solutions. The Framework should also acknowledge that the entity best positioned to leverage an individual "tool" in the toolkit is likely to vary. For instance, the entity that deploys an AI solution will likely be best positioned (as compared with the developer of the AI solution) to implement appropriate recourse mechanisms for addressing concerns that might arise through its use of the system.

Including such a conceptual distinction would be helpful to different stakeholders as they carry out risk assessments to determine the appropriate measures to adopt for AI development, deployment, and use. In addition, it would also be useful for both AI solution providers and entities that deploy and use AI to consider who the ultimate end user of the AI solution will be — in general, end-user businesses should be considered more sophisticated users than end-user individuals — and this would in turn have implications on internal risk assessments and commercial viability.

Finally, the Framework should also acknowledge the important role of government as an AI stakeholder. Government agencies should lead by example by adopting AI in their interactions with citizens[15] and investing sufficiently in infrastructure to support and deliver ethical AI solutions

[14] As reflected in the Executive Summary and section 7 "A Proposed Ethics Framework" of the Framework.

[15] For example, the Singapore government has recognized the value of governments leading by example and has set a goal for all ministries to use AI by 2023 (see the "Key Performance Indicators and Milestones" section of the Singapore government's *Digital Government Blueprint*, available at https://www.tech.gov.sg/digital-government-blueprint).

that implement the Framework's core principles. The government can also promote the development of ethical AI by fostering a policy environment that encourages controlled testing and experimentation of AI systems and by convening public-private partnerships to facilitate collaboration on AI-related issues, such as workforce development and reskilling.

BSA recommends that the Framework recognize and discuss the important roles and responsibilities that multiple stakeholders, including the Australian government, have in implementing ethical AI and managing corresponding risks.

D. Other Comments on Specific Portions of the Framework

D1. *Core Principles*

BSA recommends including further clarification on how the core principles of the Framework should be implemented. In particular, the Framework should include a greater emphasis on the point that:

- there is no "one-size-fits all solution"[16] to the emerging AI issues; and

- while the core principles serve to provide a good reference for organizations seeking to develop and use ethical AI systems, organizations are at liberty to decide when to apply (or not apply) each core principle, and the appropriate measures (e.g., processes, policies, and resources) to adopt to minimize risks identified, for specific AI systems and implementations.

D2. *Putting Principles into Practice*

BSA also has the following suggestions for improving the implementation measures in section 7.1 of the Framework:

- <u>7.1.1 (Impact assessments) and 7.1.3 (Risk assessments)</u> **—** Impact and risk assessments should go hand-in-hand for a comprehensive understanding of the tool in context. Questions around what the potential harm (impact) of the AI system would be, who is at greatest risk, how likely the risk is, and how the risk can be mitigated should all be part of the same assessment. In this regard, the nature, likelihood, and size of the harm will vary depending on the subject group in question, and different mitigation strategies may be necessary per subject group.

- <u>7.1.5 (Education, training and standards)</u> **—** This section appears to focus only on the certification of data scientists as a gateway for entering the data science profession. However, this alone does not guarantee the development of ethical AI systems. The Framework should also explore the role of general education policy to develop an ethical mindset.

- <u>Documentation</u> **—** The Framework should highlight, as part of the toolkit, the importance of maintaining appropriate documentation with respect to the implementation of the AI system in question, including for the performance of impact/risk assessments and regular monitoring of the system (as contemplated in section 7.1.7). This would aid in meeting the core principle of transparency and explainability.

[16] See the paragraph on "Implementing ethical AI" on page 8 of the Framework.

Conclusion

BSA is grateful for the Department's consultative process in developing the Framework. We hope that our comments will support the Department's efforts to promote a trusted and ethical deployment and use of AI.

Please do not hesitate to contact us if you have any questions or comments regarding our suggestions. We remain open to further discussion and look forward to further opportunities to work with the Department on the development of the Framework as well as on broader emerging technology issues in Australia.

BSA | THE SOFTWARE ALLIANCE