



Brussels, June 2020

BSA submission to the European Commission Consultation on the White Paper on Artificial Intelligence

BSA | The Software Alliance (“BSA”)¹ welcomes the opportunity to offer thoughts on the European Commission White Paper on Artificial Intelligence (“the White Paper”). 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 AI products and services. BSA members include many of the world’s leading suppliers of software, hardware, and online services to organizations of all sizes and across all industries and sectors. BSA members have made significant investments in developing innovative AI solutions for use across a range of applications. As leaders in AI development, BSA members have unique insights into both the tremendous potential that AI holds to address a variety of social challenges and the governmental policies that can best support the responsible use of AI and ensure continued innovation.

Section 1 – An Ecosystem of excellence

Software innovation is fostering the development of a range of cutting-edge technologies, such as AI, that offer great promise to improve lives and help solve intractable problems. AI solutions are already leading to improvements in healthcare, advances in education, more robust accessibility tools, stronger cybersecurity, and increased business productivity and competitiveness, impacting every sector.

AI also has the potential to generate substantial economic growth and enable governments to provide better and more responsive government services while addressing some of the most pressing societal challenges.

A flexible policy framework is necessary to enable successful deployment of AI products and services. BSA has identified five key pillars for facilitating responsible AI innovation.²

- 1) Building Confidence and Trust in AI Systems
- 2) Sound Data Innovation Policy
- 3) Cybersecurity and Privacy Protection
- 4) Research and Development
- 5) Workforce Development

¹ BSA | The Software Alliance (www.bsa.org) is the leading advocate for the global software industry before governments and in the international marketplace. Its members are among the world’s most innovative companies, creating software solutions that spark the economy and improve modern life. With headquarters in Washington, DC, and operations in more than 30 countries, BSA pioneers compliance programs that promote legal software use and advocates for public policies that foster technology innovation and drive growth in the digital economy.

BSA’s members include: Adobe, Akamai, Atlassian, Autodesk, Bentley Systems, Box, Cadence, Cloudflare, CNC/Mastercam, IBM, Informatica, Intel, Intuit, MathWorks, McAfee, Microsoft, Okta, Oracle, PTC, Salesforce, ServiceNow, Siemens Industry Software Inc., Sitecore, Slack, Splunk, Trend Micro, Trimble Solutions Corporation, Twilio, and Workday.

² For more information please visit ai.bsa.org

Section 2 – An Ecosystem of trust

BSA supports risk-based approaches to AI governance that are informed by existing law, and account for context-specific considerations in determining whether specific applications of AI should be regulated. BSA therefore welcomes the Commission's decision to adopt such an approach as the foundation for the AI White Paper.

Ensuring a balanced EU body of laws

BSA agrees with the fundamental proposition of the White Paper that the public should “expect the same level of safety and respect of their rights whether or not a product or system relies on AI.” Of course, the concerns presented by the European Commission are not unique to AI. The EU body of laws offers strong, technologically neutral safeguards against these concerns. BSA strongly recommends that the Commission take stock of this body of legislation in a targeted way, identify possible gaps and only propose new legislation if there is no other way to rectify them, AI-specific or not.

The White Paper acknowledges the challenges and promise of AI tools, and at the same time calls for a more thorough analysis of existing EU Legislation, to establish whether it is fit for purpose in protecting fundamental rights whilst fostering AI uptake. In the context of the work of the High-Level Expert Group on AI (“HLEG”), BSA prepared a detailed analysis of EU legislation impacting AI,³ which could prove helpful as the Commission moves to evaluate the sufficiency of current laws. Moreover, BSA would like to emphasize that AI is not developed in a vacuum in the EU, and that while new technologies present new challenges, the protection and enforcement of Fundamental Rights in the EU remain as strong as ever. BSA and its Members continue to work alongside EU Institutions and Member States to support a strong EU body of law that provides safeguards for fundamental rights whilst fostering innovation.

It is also important to stress that AI will be developed and deployed in an international context. If European legislation and guidelines are too prescriptive or overly rigid, AI will be developed elsewhere, and other geographies will reap the benefits of AI deployment while Europe is left behind. The international standards community is beginning to address many of the issues raised in this paper. BSA recommends that European authorities and industry fully engage in these international efforts. International engagement will be critical for ensuring that the EU approach to AI regulation is interoperable with trading partners. Such engagement has already yielded some important early successes. For instance, the Organization for Economic Cooperation and Development (OECD) Recommendation on AI represents an important first step toward establishing global norms around the governance of AI. Those norms are predicated on a risk management-based approach for enhancing the benefits of AI and safeguarding against unintended harms. The Commission can lend momentum to these positive developments by aligning future legislative efforts with the OECD's guiding principles. Moreover, to minimize the risk of international fragmentation, the Commission should consider the international regulatory landscape as it evaluates new EU legislation, and preference should be given to options that are interoperable with similar policies in foreign markets. If the Commission determines that updates to existing EU legislation are needed, the Commission should be guided by the following considerations:

³ Please refer to our submission to the HLEG on EU Legislation [here](#).

- Consistent with the risk-based, context-specific approach the Commission has endorsed, any proposed legislative changes should avoid one-size-fits-all mandates. The AI ecosystem is broad, encompassing a diverse range of technologies, use cases and wide array of stakeholders. Legislative updates must therefore be flexible enough to account for the unique considerations that may be implicated by specific uses cases. For instance, Business-to-Business (“B2B”) relations are radically different than Business-to-Consumer (“B2C”), and entail a completely different consideration and allocation of risk. In the B2B context, entities should remain free to use contractual negotiations as a mechanism for allocating risks, liabilities, and obligations in a manner that corresponds to the nature of the transaction.
- To the extent new statutory obligations are contemplated, they should account for the unique roles and capabilities of the entities that may be involved in an AI system’s supply chain. To that end, the Commission should ensure that any new regulatory obligations (and associated liabilities) fall on the entity that is best positioned to both identify and efficiently mitigate the risk of harm that gave rise to the need for a regulation. In many circumstances, only the entity that has deployed an AI system will be in a position to monitor whether it is operating as intended and intervene when necessary to mitigate risk.
- Defining AI properly will also be crucial. The term “artificial intelligence” can be, and often is, used to describe a vast array of technologies. These technologies, in turn, can be used in a nearly infinite range of scenarios. Consistent with the Commission’s goal of carefully focusing regulatory mandates on high-risk scenarios, it will be important to ensure that the definition of AI is not so broad as to sweep in thousands of everyday products and services.

A systematic risk-based approach to AI

BSA agrees that future legislative proposals should focus on high-risk scenarios where the deployment of AI-based technologies poses a threat to fundamental rights. 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. To this end, it will be important to carefully assess scenarios that should be deemed as high-risk and hence be subject to legal requirements. BSA strongly recommends ensuring stakeholder involvement in this context as much as possible, especially as it will be sector-dependent as much as use-case dependent. BSA and its Members have participated, and intend to continue to be active participants, to EU and Member States stakeholder consultations. BSA Members are uniquely positioned to provide essential insights in the assessment of high-risk scenarios and use-cases of AI.

BSA welcomes the decision to take an incremental approach by limiting regulation to AI systems that are (1) deployed in a high risk sector and (2) used in a manner that significant risks are likely to arise. Moreover, due consideration should be given to AI applications that enhance human decision-making, whereby the risk consideration – even when the two above conditions are fulfilled – is inevitably balanced by the human involvement and control. Furthermore, BSA urges the Commission to extend this two-pronged approach to all possible high-risk scenarios, rather than identifying specific sectors where – regardless of its purpose and use – AI would be considered high-risk by default. This would allow for a more homogeneous application and understanding of the possible requirements for high-risk AI, providing for the necessary proportionality and legal certainty as AI technologies and tools are developed and deployed.

Ensuring that the definition is appropriately tailored will be critical. Given the potentially far-reaching requirements of new legislative requirements for high-risk AI, it is crucial for AI developers and users to be able to determine with certainty if their application might fall within the scope of high-risk. The complexity of defining “high-risk” AI is exacerbated by that fact that AI may be developed for a multitude of uses – and determining whether it is “high-risk” will turn on how it is deployed (i.e., whether it is deployed by an end-user in a high-risk sector and used in a manner that creates a significantly likelihood of risk.). The methodology behind the definition of high-risk sectors needs to be precise and robust, with only limited exemptions. This will guarantee that the list remains targeted and up to date as new technologies and use cases emerge.

BSA agrees with the Commission’s analysis that legal requirements for high-risk AI applications “should be addressed to the actor(s) who is (are) best placed to address any potential risks” (White Paper p. 22). In many cases – especially in the cases of general-purpose AI systems – developers will not be in the position to know whether the technology is being deployed by an end-user in a manner that meets the definition of high-risk. Similarly, in B2B relations the allocation of risk will be one part of the contractual agreement between two entities, and once again the developer will not be best place to establish whether the application is to be deployed in an high-risk scenario, and what obligations that may entail in the specific sector. Developers are better placed to describe the capabilities and limitations of an AI system, while disclosing the way AI can impact the fact of AI use to people likely to be affected by it will typically need to be the responsibility of the deployer.

In this context, and in a similar vein in the liability context (please see below in the relevant section for more information), the Commission may want to draw from existing concepts for establishing which entity is “best placed to address any potential risk”, i.e. the entity that determines the purpose of the AI, similar to the concept of a “controller” under the GDPR. Article 29 Working Party guidance on controllers and processors (WP 169) describes this party as the “determining body” that decides the “how” and the “why” of a processing operation. Applying this concept in the context of AI, the “AI controller” will generally be the deployer of an AI system (e.g., a vehicle manufacturer that integrates an AI-driven language recognition system into an automobile, or a bank that uses an AI tool to score consumers for loans). In some instances, it may also be the operator of the AI system (e.g., a physician using assistive tools during surgery).

Under the GDPR, controllers and processors have different levels of responsibility for achieving privacy outcomes that reflect their different roles. In particular, controllers have primary responsibility for satisfying certain legal privacy and security obligations and for honoring data subject rights requests. On the other hand, processors, which handle data on behalf of the controller to implement the controller’s objectives, are responsible for securing the personal data they maintain and following the instructions of a controller, pursuant to their agreements with relevant controllers. The processor/controller distinction not only provides organizations with a clear picture of their respective legal obligations, it also helps to ensure that data subjects rights are adequately protected.

This key distinction could also help inform the Commission’s AI workstreams, which will have to focus on sectors with very different definitions and approaches to risk management. In the context of enterprise AI, the tools that BSA companies provide are generally AI systems that facilitate human decision-making, without replacing human decision-making. With this in mind, it becomes clear that a company using an AI service to enable its employees to make a decision acts as a controller in deciding how and why that data is processed, and the AI system is used as a tool for processing data on behalf of that controller. Accordingly, the company developing and providing the AI tool is appropriately treated as a processor.

Risk assessment for AI technologies and tools

BSA recommends to the Commission not to establish pre-marketing conformity assessment for AI systems, as such obligations are liable to turn into barriers to enter the market. A more scalable approach would be self-attestation, which would also be least likely to unnecessarily extend time to market or unduly burden smaller operators. BSA believe that prescriptive regulation of AI, requiring for example that every possible use of an AI system is “fair” or “unbiased”, will likely be unworkable in practice.

BSA urges the Commission not to pursue a regulatory scheme based on prescriptive conformity assessment requirements. The risks that AI poses and the appropriate mechanisms for mitigating those risks are largely context-specific. The appropriate mechanisms and standards for training data, record keeping, transparency, accuracy, and human oversight will vary depending on the nature of the AI system and the setting in which it is being deployed. The Commission should therefore avoid creating prescriptive, one-size-fits-all requirements around these categories. Such ex-ante requirements could impede efforts to address the very risks they are intended to address, add unnecessary costs and require extremely complex compliance checks.

Given the nascent nature of the technology and sociotechnical quality of many of its most significant challenges, BSA believes that a governance-based approach to legislation, which identifies broad objectives and the processes that developers and deployers should follow to achieve them, will be more effective than a prescriptive one.

Consistent with a governance-based approach, the Commission should articulate a framework that will enable stakeholders to perform an “impact assessment” on high-risk AI systems. In this context, BSA recommends building upon the work done by the HLEG and many AI developers on the Assessment List, in particular as it may constitute a template for future assessment tools. The goal of these governance processes should be to help developers and deployers of covered AI systems identify and quantify any relevant risks of harm to individuals or society and, where those risks are determined to be significant, to implement measures to mitigate against them. Importantly, impact assessments allow for a more context-specific evaluation of the types of risk mitigation measures that are available, and which is ideally suited for the particular deployment scenario.

BSA acknowledges the Commission’s concern that the deployment of biometric identification systems can implicate heightened risks for fundamental rights. BSA welcomes the White Paper’s recommendation for the Commission to launch an inquiry to examine the appropriate regulatory framework for biometric systems. While EU law already provides clear parameters for assessing the lawfulness of biometric technologies from a data protection perspective, the rules that govern the ethics and other risks of Facial Recognition Technologies (“FRT”) deployments are less well defined. For that reason, the Commission should consider specific rules governing the use of FRT by the public sector in particular, given the heightened risks inherent in governmental use of this technology.

BSA agrees that public trust in AI is essential for “[promoting] the overall uptake of the technology”. However, we would urge the Commission not to pursue the creation of a blanket voluntary labeling system for all no-high risk systems. Given the diverse range of AI products and services that will be considered “no-high risk AI applications”, a one-size-fits-all labeling scheme would be unworkable. The benchmarks for evaluating whether AI systems are trustworthy are likely to be highly variable, driven in large part by system functionality and deployment context. The relevant benchmarks for evaluating the trustworthiness of an AI system that recommends restaurants are likely to be quite different from those that are relevant to an AI system that is designed to identify what objects are in a photograph. A methodology for a labelling system that

could apply to the entire universe of “no-high risk AI” would necessarily be very complex, which would limit customers’ understanding and engagement. Similarly, the governance of such a scheme would be exceedingly complex and would necessarily have to cover diverse sectors and technologies – likely in almost all industrial EU activities.

A strong stakeholder engagement structure

BSA commends the Commission’s intention to maintain a strong focus on the governance of future AI legislation and rules, especially in the implementation and enforcement phase. In this context, as the Commission correctly notes, stakeholder engagement systems already exist in certain sectors. BSA recommends ensuring that clear language for broad stakeholder involvement is included in future legislation, to promote a beneficial interaction between AI developers - which may not have extensive experience or presence in a specific sector now deemed high-risk - and deployers. As legislation is implemented and enforced, BSA urges the Commission to retain a coordinating competence for stakeholder engagement throughout the legislative process, and especially in the implementation and enforcement phase.

In particular with regards to enforcement, the Commission should continue to endeavor to empower the existing sectoral authorities – for “high risk” sectors – which will be best placed to provide fundamental insights as to how AI applications and legislation may impact the specific sector. At the same time, the Commission should leverage its existing coordinating role to ensure that enforcement authorities remain in contact and apply harmonized principles, as well as continuous work with practitioners to update best practices and guidelines.

BSA recommends ensuring that the EU continues to work alongside global partners – in the private and public sector – so as to incorporate international standards and best practices in the European workstream on AI regulation. Governance of AI in the EU should be done in a way that prevents unnecessary barriers to transatlantic trade and investments. Dialogue with the US should ultimately lead to the development of de facto global standards for AI governance, based on common democratic values.

Section 3 – Safety and liability implications of AI, IoT and robotics

BSA believes that any Commission effort in the space of updating EU liability rules, should be guided by the following principles:

- **Liability rules should be technology neutral.** Products should not be subject to unique or heightened liability rules simply because they integrate AI. That approach will deter producers from offering AI-driven products that may in fact make consumers safer out of fear that doing so could expose them to increased liability. Further, because AI is deployed in a virtually limitless set of scenarios — from AI-powered tangible products and fully autonomous systems to AI software that merely assists or informs human action and decision-making, and including systems that evolve through machine learning and / or can be personalized — a unique liability regime for AI would be all but impossible to implement. Given the broad nature of “AI” as a category of technology, it is difficult to conceive of a coherent “one-size-fits-all” approach for AI liability that would make sense in the context of both driverless cars and, for example, AI systems that help businesses identify customer preferences and adjust their service offerings to them. Instead, the liability rules applicable to AI should be technology-neutral while allowing courts to take

full account of the differences between AI systems, the relative risks they pose, and the contexts in which they are used (just as EU product liability rules do today).

- **Any changes to the EU’s existing liability regime should be driven by a clear and demonstrated need.** The EU current product liability regime, set out in the Product Liability Directive (“PLD”), has worked well in a wide variety of contexts for over 30 years. The PLD sets out clear, well understood and time-tested rules that apply across a wide range of products, including those with embedded software. The Directive is complemented by national civil liability frameworks that reflect Member State legal traditions and principles, including evidentiary rules, damages regimes, and contract and tort law. Although it is true that the EU liability regime has not been updated in some time, that is because the available evidence suggests that it is working well. It is also important to underline that consumers have the possibility to obtain compensation for possible harms due to AI, or other products or services under the current regime. Rushing to change liability rules without further focused assessment of the status quo, risks chilling innovation into socially beneficial (and even safety-enhancing) AI, with little positive benefit for consumers.
- **Liability rules must align with the safety regime.** Product safety and product liability are complementary regimes designed to pursue complementary policy goals: to encourage the development of safe products, and to ensure that consumers have recourse when things go wrong. Product safety regimes achieve the first goal by requiring producers to take steps to ensure that the products they market are safe; liability rules achieve the second goal by allowing consumers to recover for the harms caused by defective products. In the case of high-risk AI, the EU is currently considering a wholly new safety framework (e.g., to require data quality, transparency, robustness). The final shape of that new safety framework could be highly relevant for the corresponding liability rules (e.g., if high-risk AI products comply with applicable standards, this arguably should provide a defense to liability). Any changes to the liability framework should await the conclusion of this exercise to ensure the regimes are appropriately aligned.