# AI for Europe

Software innovation is fostering the development of a range of cutting-edge technologies, such as artificial intelligence (AI), that offer great promise to improve lives and help solve intractable problems. AI is 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.[1]

At the same time, AI that is not developed, trained, and deployed responsibly does carry the risk of significant negative consequences that can result in an erosion of public trust in AI. BSA therefore supports industry efforts to provide users of AI systems with the information necessary to instill confidence that such systems were developed responsibly and are operating as intended. Facilitating increased understanding and promoting trust in the use of AI technologies is an important priority.

Facilitating increased understanding and promoting trust in the use of AI technologies is an important priority.

We likewise support a flexible policy framework that encourages responsible AI practices that are critical to the successful deployment of AI products and services. BSA highlights the below recommendations that would help in this endeavor:

**Evidence-Based Approach.** Carry out an in-depth study of how the already strong and effective body of EU law applies to AI, before considering new legislation.

**Risk-Based Framework.** Build a two-tiered risk-based approach that is sector and use-case specific.

**Balance the Allocation of Responsibility.** Ensure that responsibility for risks is assigned to the actor best placed to identify and mitigate potential harms before they arise. To that end, the controller/processor distinction in GDPR could be adapted to AI deployers and AI developers.

[1] BSA Response & Recovery Agenda, May 27, 2020, https://www.bsa.org/policy-filings/bsa-response-recovery-agenda

## Evidence-Based Approach

A fundamental proposition of the European Commission's AI White Paper is 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." Developed and strengthened over the years, the EU body of laws offers strong, technology-neutral protections that address multiple concerns pertaining to AI. 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,[2] which provides an overview of how EU law already responds to many of the challenges posed by new technologies. BSA recommends that the European Commission takes 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.

Moreover, while new technologies may 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 while fostering innovation.

It is also important to stress that AI will be developed and deployed in an international context, and the international standards community is beginning to address many of the concerns raised around AI. 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. Furthermore, to minimize the risk of international fragmentation, the European 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.

While new technologies may present new challenges, the protection and enforcement of Fundamental Rights in the EU remains as strong as ever.

## Risk-Based Framework

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, BSA supports the European Commission's approach of limiting regulation to AI systems that are (1) deployed in a high-risk sector and (2) used in a manner that gives rise to significant risks. BSA cautions against classifying certain sectors as per se high-risk, whereby any AI tool deployed would be considered high-risk regardless of its purpose and use. Applying the abovementioned two-pronged approach would much better respond to the concerns the Commission seeks to address. It would also 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.

As part of this framework, **the extent of human-in-the-loop involvement should be considered.** In such cases, AI applications may be used to enhance human decision-making, and the risk consideration—even when the two above conditions are fulfilled—is inevitably mitigated to some extent by the human involvement and control.

[2] BSA Letter to the HLEG, June 6, 2019, https://www.bsa.org/files/policy-filings/06062019bsasubmissionaihleg.pdf

The touchstone of this approach should be the risk posed by specific uses of an AI technology. Given the nascent nature of AI technology and sociotechnical quality of many of its most significant challenges, a governance-based approach to legislation, which identifies objectives and the processes that developers and deployers should follow to achieve them, would be more effective than a prescriptive one. This is especially true because the EU body of laws already provides for strong safeguards for consumers and businesses.

**Any proposed new legislative instruments should avoid one-size-fits-all mandates.** The AI ecosystem is broad because it includes a diverse range of technologies and use cases, and involves a wide array of stakeholders. 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. These categories do not lend themselves well to prescriptive and one-size-fits-all requirements. 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.

Consistent with a governance-based approach, BSA recommends articulating a framework that will enable stakeholders to perform an "impact assessment" on high-risk AI systems, building upon the work done by the HLEG and many AI developers on the Assessment List for Trustworthy Artificial Intelligence. The goal of these governance processes should be to help developers and deployers of covered AI systems document the processes by which they have identified and quantified any relevant risks of harm to individuals or society, as well as the measures they have taken to mitigate such risks. Importantly, impact assessments allow for a more context-specific evaluation of the types of risk mitigation measures that are available, and which are best suited for the particular deployment scenario. **A combination of strong stakeholder engagement in designing best practices for risk-assessment, and legislation that is built upon such system, is more likely to integrate and encourage innovation within clear legal parameters and requirements.**

## Balance the Allocation of Responsibility

The European Commission's conclusion—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"—should be the guiding light in establishing how risk management and liabilities are allocated. 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. To the extent new legislation is contemplated, it should account for the unique roles and capabilities of the entities that may be involved in an AI system's supply chain. **Any new regulatory obligations (and associated liabilities) should 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.**

Legislative updates must be flexible enough to account for the unique considerations that may be implicated by specific uses cases. Business-to-Business (B2B) relations are not the same as Business-to-Consumer (B2C) uses, and therefore entail different considerations and allocation of responsibility and 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.** In B2B relations, the allocation of risk and responsibility will be one part of the contractual agreement between two entities, and that allocation should be based on which party is in the best position to establish safeguards and mitigate the risk of harm.

Business-to-Business relations are not the same as Business-to-Consumer uses, and therefore entail different considerations and allocation of responsibility and risk.

Existing EU legislation may serve well in helping establish which entity is "best placed to address any potential risk[]". The entity that determines the purpose of the AI is often similar to the concept of a "controller" under the GDPR.[3] 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), **the "AI processor" will generally be the developer of the AI system** (e.g., the entity that developed all or part of the AI-driven language recognition system as per the example above).

This key distinction could also help inform different AI workstreams, focusing on sectors with very different definitions and approaches to risk management. **Developers are often better placed to describe the capabilities and limitations of an AI system, whereas the performance of a context-specific impact assessment and disclosing the fact of AI use to people likely to be affected by it will typically need to be the responsibility of the deployer.**

Under the GDPR, controllers and processors have different responsibilities 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. It is nevertheless important to stress that adapting the controller/processor distinction to the developer/deployer relation in AI will still need the necessary nuance. The context and purpose of AI tools should remain a key guiding principle also in this context.

Adapting the controller/processor distinction to the developer/deployer relation in AI will still need the necessary nuance. Context and purpose of AI tools should remain a key guiding principle.

Beyond the developer/deployer concepts, it is equally important to note how risk consideration vary greatly depending on the sector. In particular, in B2B relations risk is often allocated on a contractual basis. Due to the complex and often diverse layers in a supply chain, contractual agreements are often the preferred avenue to regulate relations between businesses. **BSA recommends including language allowing companies to allocate risk and responsibilities on a contractual basis in the B2B space.** This would be beneficial both for existing B2B contracts and ultimately for protecting consumers from potential harm caused by the misuse of AI.

BSA commends the European Commission's intention to maintain a strong focus on the governance of future AI legislation and rules, especially in the implementation and enforcement phase. **BSA recommends ensuring that clear language for broad stakeholder involvement is included in future legislation, to promote a beneficial interaction between AI developers and deployers.** As legislation is implemented and enforced, the European Commission should retain a coordinating competence for stakeholder engagement throughout the legislative process, and especially in the implementation and enforcement phase.

[3] 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.