Introduction: Rationale and Foresight of the Guidelines

BSA | The Software Alliance[1] welcomes the opportunity to offer thoughts on the High Level Expert Group (HLEG) Draft Guidelines. 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.

The formation of the HLEG is a unique opportunity for Europe's leading experts from industry, academia, and civil society to help the European Commission develop a "coordinated approach to make the most of the opportunities offered by AI and to address the new challenges that it brings." We agree with the Commission that the success of such a framework will turn in large part on whether it fosters an "environment of trust and accountability around the development and use of AI."

The HLEG requested comments on the potential concerns raised by AI. BSA agrees that it is important to consider potential short and long term concerns. BSA supports efforts to deploy AI responsibly. We urge the HLEG to be cautious in recommending policy action now based on speculative concerns, however. The Draft Guidelines make the important point that "it should be noted that no legal vacuum currently exists, as Europe already has regulation in place that applies to AI". BSA believes that building trust in Artificial Intelligence system is one of five key pillars[2] that we have identified for the development of Responsible AI. These pillars are:

1) Building Confidence and Trust in AI Systems: please refer to the list below (Chapter I) of five principles for more information on this pillar.

---

[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 60 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, Apple, Autodesk, Bentley Systems, Box, Cadence, CNC/Mastercam, DataStax, DocuSign, IBM, Informatica, Intel, Intuit, MathWorks, McAfee, Microsoft, Okta, Oracle, PTC, Salesforce, Siemens PLM Software, Slack, Splunk, Symantec, Trend Micro, Trimble Solutions Corporation, Twilio, and Workday.

[2] *For more information about these pillars, please visit ai.bsa.org*

Avenue des Arts 44
1040 Brussels
Belgium

P +32 (0)2 274 13 10
W bsa.org
EU Register of Interest Representatives 75039383277-48

2) Sound Data Innovation Policy: The exponential increase in data has fueled advances in machine learning and AI. Facilitating the development of AI requires sound data innovation policies.
3) Cybersecurity and Privacy Protection: BSA advocates 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: Investment in education, research, and technological development will be integral to continued development of AI technologies and global economic growth.
5) Workforce Development: The increasing use of and demand for technology is creating new types of jobs, in every sector of the economy, that require an evolving set of skills.

The discussion fostered by the HLEG around Trustworthy AI is a fundamental step in ensuring that the EU has a leadership role in innovative technologies. At the same time, excessively prescriptive guidelines can be counterproductive in a field that is in rapid development, such as AI.

The HLEG's work is an important step in initiating a larger debate across the EU on Artificial Intelligence, and BSA supports its ethical framing. BSA recommends establishing a clear process and timeline to update the guidelines with stakeholder involvement to ensure the HLEG recommendations stay relevant and in line with technological development. While the AI Alliance is certainly an effective platform for engagement, BSA recommends the creation of additional fora at the EU level to ensure stakeholders involvement, as well as establishing a timeline for a short-term evaluation of the guidelines.

BSA welcomes the Draft Guidelines as an excellent effort to strengthen the EU's role as a global leader in ethical and responsible development of innovative technologies. Overall, the Draft Guidelines seek to strike the right balance between complex ethical dilemmas and the need to support AI development in the EU. In their current formulation, they are not merely a compilation of values and principles, rather they also provide guidance on how to implement these principles. BSA considers the Draft Guidelines an excellent step and foundation for the AI conversation at the EU and global level. Moreover, the HLEG's recognition that Trustworthy AI has two components – ethical purpose and technical robustness - is an important guiding principle for any work done on AI around the world.

At the same time, BSA also recommends that the HLEG takes into consideration similar international efforts in the space of AI and ethics, to ensure that the EU can contribute to the global discussion and development of AI tools. An international perspective on how to realize trustworthy AI, and assess it, would ensure that the EU remains competitive on the global markets, whilst contributing to strengthen trust in new technologies.

On a more general point, BSA recommends a more positive approach to Artificial Intelligence. The HLEG correctly noted the tremendous potential AI has to spur economic growth across every industry sector, improve human decision-making in ways that will make the world more inclusive, and enable cutting-edge breakthroughs on vexing social challenges such as climate change and cancer research.

BSA therefore recommends:

- Creating additional instruments for meaningful and routine stakeholder consultations after the publication of the final guidelines;
- Strengthening the role of the members of the AI Alliance platform and create additional means of stakeholder involvement;
- Establishing short-term and long-term timelines for evaluation of the guidelines;
- Establishing a clear process to amend and update the guidelines to ensure they remain relevant and in line with technological development;
- Ensuring that the guidelines are informed by and contribute to international efforts.

Finally, we note the new definition of AI provided by the HLEG.

BSA appreciates the Guidelines' thoughtfulness in proposing a possible definition of AI. We note, however, that many solutions in use today that are described as having an AI component make connections, reveal correlations, or provide other insights that humans then use to decide on a course of action, but do not necessarily decide "the best action(s) to take" as stressed in the Guidelines' definition. BSA is still considering the implications of the proposed definition and may provide further input on it after conducting a more in-depth analysis.

Turning to the specific Chapters developed by the HLEG, BSA respectfully submits the following comments for consideration:

Chapter I: Respecting Fundamental Rights, Principles and Values - Ethical Purpose

BSA supports discussions of high-level ethical principles pertaining to Artificial Intelligence development, especially within the frameworks of the EU Treaties and Charter of Fundamental Rights. While such an exercise is important to begin a principles-oriented debate at the European level, it is also important to underline that AI development and use is happening globally. The EU would greatly benefit from considering the existing best practices, developed by public and private sector, at the global level. It is important both that EU values and priorities are fully taken into account at the international level, and that international standards and principles are incorporated in any European effort.

While the Guidelines will need to reflect European values, we must also keep in mind that AI will be developed and deployed in an international context. If European good practice or European guidelines are too draconian, prescriptive or overly rigid, AI will be developed elsewhere and other geographies will reap the benefits of AI innovation while Europe is left behind. The international standards community is beginning to address many of the issues raised in the Guidelines. BSA recommends that European authorities and industry fully engage in these international efforts.

Avenue des Arts 44
1040 Brussels
Belgium

P +32 (0)2 274 13 10
W bsa.org
EU Register of Interest Representatives 75039383277-48

As a global organization, BSA has also developed a set of 5 principles for building confidence and trust in AI systems, which are consistent with the 5 principles the HLEG suggests in the Draft Guidelines. The BSA principles are:

a) **Fairness**—considering measures to evaluate AI systems to help recognize improper or unconscious bias;

b) **Accuracy**—acknowledging the importance of data quality and, where feasible, identifying sources of error in data inputs and system outputs;

c) **Data Provenance**—considering measures that could facilitate evaluation and documentation of data used to train AI systems, how those data are collected, and how data is used over time within AI systems, consistent with any other data retention obligations;

d) **Explainability**—exploring how to provide reasonable explanations of how AI systems operate; and

e) **Responsibility**—considering whether processes are available to address unexpected issues that may arise after AI products and services are deployed.

At a high level, both the Ethical Principles discussed in the Draft Guidelines and the BSA principles stress the importance of designing and using AI as an understandable tool to aid human decision making and improve economy and society.

BSA welcomes the HLEG recognition of the diversity of AI applications, and the importance in avoiding the creation of a "one-size-fits-all" regime. We believe that contextual considerations merit greater attention in the Guidelines. The degree of risk of individual or societal harm, and the potential severity of such harm, will vary enormously depending on the specific AI application at issue. In fact, many of the ethical issues identified in the Guidelines only arise for AI systems that have a consequential – or meaningful – impact on individuals. BSA therefore urges the HLEG to make clear at the outset of the Guidelines that the recommendations should be tailored to each specific implementation of AI depending on a careful and thorough risk assessment.

Engagement at the principles-level is an important step in strengthening trust in AI tools. To that end, BSA commends the HLEG's efforts to ensure that the Draft Guidelines take into account the multitude of diverse applications of AI, as well as the technical considerations related to enacting ethical principles in this space.

BSA encourages the HLEG to recognize more explicitly that AI policy involves trade-offs, and therefore a risk-based approach, tailored to the circumstances, will be necessary. BSA also recommends that the Guidelines adopt a broad understanding of beneficence. In fact, AI can be a tool to improve wellbeing, but it can also serve more neutral objectives whose direct individual or social benefits are less clear.

Chapter II: Realising Trustworthy AI

Trustworthy AI is a critical objective and BSA supports the HLEG efforts to provide a general framework for practitioners to achieve it. As the HLEG has rightly noted the "Guidelines are not meant to stifle AI innovation in Europe, but instead aim to use ethics as inspiration to develop a unique brand of AI". We therefore recommend a light-touch approach on many aspects of the Guidelines, in the interest of working with practitioners around the world to develop a framework that supports innovation while providing workable guiding principles. To that purpose, BSA suggests streamlining the below list of 10 requirements, and aggregating some of the requirements to ensure that the list is of easier use.

In particular, we would like to provide the following comments to the 10 requirements the HLEG identified in the Guidelines:

1) Accountability: BSA agrees accountability in development is needed. We caution that this is among the areas in which the mechanism has to be sufficiently flexible to accommodate different use cases and means of deployment.

2) Data governance: BSA believes that a trustworthy AI should respect the principles of Accuracy (i.e. acknowledging data quality and, where feasibly, identifying sources of error in data inputs and system outputs), Data Provenance (i.e. considering measures that could facilitate evaluation and documentation of data used to train AI systems, how those data are collected, and how data is used over time within AI systems, consistent with any other data retention obligations) and Fairness (i.e. considering measures to evaluate AI systems to help recognize improper or unconscious bias). The Guidelines' concept of "data governance" should be broader and reflect the fact that governance structures necessary to develop AI ethically include a broad range of engineering and design practices as well (e.g. access controls, systems documentation), BSA therefore urges the HLEG to recognize that data governance is complex in practice and will need to be tailored to individual scenarios.

3) Design for all: BSA supports broad access to AI products, in particular as many applications of AI will greatly benefit underserved portions of the population. To that end, we would suggest including recommendations for public sector support to industry to develop products with high accessibility. Nevertheless, we would also caution against overly prescriptive requirements in the development and design phase, as flexibility in innovation is an integral part of any development process.

4) Governance of AI autonomy (Human oversight): BSA agrees that human oversight is an important principle, given the diversity of AI tools, and the different technical considerations they would entail (e.g. in machine-to-machine applications), BSA would suggest to also consider context and purpose of an AI technology with this requirement.

5) Non-discrimination: limiting bias and unfair discrimination are fundamental objectives. BSA believes that trustworthy AI should respect the principle of Fairness (i.e. considering measures to evaluate AI systems to help recognize improper or unconscious bias) and has as well put forward a number of best practices recommendations to limit the effects of unfair bias in AI development

Avenue des Arts 44
1040 Brussels
Belgium

P +32 (0)2 274 13 10
W bsa.org
EU Register of Interest Representatives 75039383277-48

(please refer to our comments on Chapter III). At the same time, BSA would like to stress that measures in place to limit bias should not be considered absolutely foolproof. In some instances, it may be necessary and/or intended to consider certain individual characteristics (e.g. in the healthcare sector, diagnosis tools might need to consider age, sex or personal background as factors for diagnosis, as they might lead to higher propensity for some diseases or different reactions to cures). BSA would stress the importance of highlighting the need to protect against unfair discrimination and bias.

6) Respect for (& Enhancement of) Human Autonomy: BSA agrees that fundamental and constitutional rights need to be safeguarded with the progressive deployment of AI technologies. While the principles defined in Chapter I are designed to improve and guide AI development, it is important to stress that AI will function in an already strong rule of law system within the EU, and will be designed to respect and strengthen that system.

7) Respect for privacy: BSA fully agrees that trust and privacy are foundational to the development and adoption of AI. Beyond the necessary full compliance with GDPR, BSA promotes best practices globally that increase the transparency of personal data collection and use; enable and respect informed choices by providing governance over that collection and use; provide consumers with control over their personal data; provide robust security; and promote the use of data for legitimate business purposes.

8) Robustness: BSA is a strong advocate of data accuracy, resilience and cybersecurity. The more complex a system, the more important these principles and practices become. As mentioned above, BSA has developed a set of principles and best practices to ensure trust in AI tools, and in particular strongly believes that Accuracy, Data Provenance, and Responsibility should be the guiding principles for robust and trustworthy AI systems. With regards to cybersecurity, BSA has developed a wealth of materials to promote cybersecurity awareness, while protecting privacy and safety (for more on this, please refer to bsacybersecurity.bsa.org).

9) Safety: BSA fully supports safe systems and believes that trust can only be earned through safety in practice. Safety is a fundamental component of trust in AI tools, and BSA is fully committed to the highest standards in AI development and deployment.

10) Transparency: Explainability is a key principle to ensure trustworthy AI. In particular, BSA believes that explainability (which is a more accurate term to use in providing an explanation of the AI system's approach in useful terms for the user) will inevitably vary due to context and purpose, it is also important to develop strong best practices and support information for all users of AI tools. It is as well fundamental to acknowledge that any effort in the field of transparency will need to take into account the developers' ability to innovate and provide cutting-edge services. BSA believes that these efforts would be better developed if led by industry and developers as AI tools are deployed. In addition, achieving transparency can be complex and highly dependent on a host of variables, precluding a "one-size-fits-all" approach.

When it comes to technical methods to achieve Trustworthy AI, and particularly with regards to traceability and auditability, BSA believes that the nature of auditability will be heavily context dependent. In complex scenarios, third party auditors and expert controls will be more effective for technical support. In still other scenarios, internal organizational auditing and controls may suffice. In light of this, the Guidelines should do more to acknowledge that effective auditing, depending on the context, can include any of those mechanics.

BSA is a strong supporter of efforts to create risk-based regimes that support solutions to significant technical and operational issues in new technologies. The HLEG should consider advancing the use of risk assessments for AI as a tool for companies to interpret how technological, operational, and policy controls, requirements and standards can support implementation of Trustworthy AI. Guidance should also embody a risk-based approach to deciding where companies should most effectively focus their efforts.

Risk management is at the core of advancing trustworthy AI. Those developing or implementing AI should be responsible for conducting appropriate, robust risk assessments. Identified risks should be mitigated through effective safeguards, such that the benefits of the AI implementation outweigh the residual risks. This is fundamental to ensuring that users and other stakeholders are protected and safeguarded.

Chapter III: Assessing Trustworthy AI

BSA strongly supports any effort by the HLEG, and subsequently by the Commission, to involve a broad audience of practitioners in developing best practices and guidance for assessment instruments. As the HLEG rightly notes in the Draft Guidelines, any assessment instrument, tool or best practice will have to be considered in context, and given the specific purpose of each AI application. We also support the HLEG's efforts to provide such guidance, which is important to enable innovators to understand in a practical way how ethical principles can be deployed.

As a global organization whose members are at the forefront of AI development, BSA has developed a list of practices for responsible AI deployment. These practices, provided below, are necessarily high-level because of the numerous use cases and deployment models:

- Conducting in-house testing and evaluation of AI systems to ensure they meet their specified goals;
- Developing guidelines and providing necessary resources to developers to help evaluate fairness and guard against improper bias;
- Identifying persons with relevant expertise who are responsible for addressing significant problems identified with operating AI systems;
- Ensuring subject matter experts, especially those with knowledge of the policy landscape in which the AI system will be deployed, are available to assist computational scientists in the design and implementation phases;

Avenue des Arts 44
1040 Brussels
Belgium

P +32 (0)2 274 13 10
W bsa.org
EU Register of Interest Representatives 75039383277-48

- Providing descriptions of procedures used to assess the quality of data inputs and address errors identified in outputs;
- Providing general descriptions, where appropriate, of training datasets that AI systems use to learn;
- Developing mechanisms for consumers to request information, obtain guidance and address potential concerns;
- Continuing monitoring after product release to detect and address unintended outcomes;
- Providing visual aids and/or plain language explanations that communicate important facts about AI systems and their operation; and/or
- Supporting continued research and analysis of transparent modeling.

The HLEG asks specifically about how an assessment would work in four "use cases": (1) healthcare diagnose and treatment; (2) autonomous driving/moving; (3) insurance premiums; and (4) profiling and law enforcement. BSA agrees with the draft Guidelines that precise questions relevant to assessment of any particular AI system will vary depending on the use case. The difficulty is that these use cases are themselves broad categories. There are numerous different uses and deployment models for AI within each category, that have different levels of risk based on the nature of data sets, the time for human intervention, and numerous other factors. For these reasons, BSA developed the more general list of best practices above, and urges the HLEG to consider a similar, flexible approach.

Yours faithfully,

Thomas Boué

Director General, Policy – EMEA
BSA | The Software Alliance