Joint Position on AI White Paper

We, the undersigned trade associations and AI Platforms from Central and Eastern Europe are convinced that Europe has the expertise to fully open unprecedented opportunity which comes with the development of artificial intelligence. **Europe is well positioned to play a leading role in AI research and application**. The Commission has rightly identified a need for focus on investment in and deployment of AI to ensure its benefits across the board continue to grow, in particular in the context of the future economic recovery. In terms of immediate priorities, the Commission should centre its efforts on:

- *supporting the research and innovation community* - the proposal to create a lighthouse centre for AI research, innovation and expertise in Europe is an excellent initiative, although it is vital to strike the right balance in the degree of centralisation
- *ensuring the right skills are in place* so all are able to prosper from the AI benefits and that sustainable development is supported by a pipeline of population that possess advanced digital skills
- *supporting SMEs* to help companies understand and exploit the potential benefit of AI for their business, and
- *promoting public sector adoption of AI*.

Supporting human capital

Supporting educational curricula, and other measures will not be sufficient, given the current demand of European industry and research organizations for AI specialists. Much more support for talent identification, development and acquisition must be acknowledged. New AI work-force must undoubtedly be also acquired on the international level. The document should mention ways to bring AI workforce from outside Europe - for example from Ukraine or India, where IITs and IIITs are educating excellent AI students that are willing to come to Europe (at this point, we are benefiting from the US immigration restrictions and have a unique window of opportunity). The document mentions acquiring talent only for the "lighthouse centre of research and innovation" (section C. Skills) while it should be a general principle and process aimed at all players across European AI (R&D institutions, industry). The procedure to get AI specialists to work in Europe should be as fast and straightforward as possible.

Cooperation on future development of AI

The White Paper lacks a concrete strategy of cooperation with big global technological companies. "Global players" are mentioned in section H but with no details. Data is essential for any AI research, development and deployment, and these companies hold most of European end-user data, and probably a big part of business data. They also have top AI R&D groups deeply involved in basic and applied AI research. It is crucial for future development of AI in the EU to present a strategy or plan on how Europe can cooperate with these companies on mutually beneficial terms.

Europe must realize that the path to the further development of is through developing European capabilities and continue cooperation with technology leaders around the world. Most countries where the rule of law and the democratic establishment exist, share concerns about data protection, cyber security and future industrial capabilities. The Commission should seek cooperation with credible international partners such as OECD countries who can guarantee to comply with European rules and regulations including GDPR.

Language technologies

The position of language technologies is not adequate in the White Paper - despite significant amount of research and successful industrial companies in this field, the document actually contains no mention of machine translation and natural language processing, speech recognition is mentioned only in a footnote. The variety of languages (24 official ones) in the EU distinguishes Europe from "linguistically homogeneous" regions such as the US or China. The digital single market will not work without serious investments into the R&D of speech and language technologies (voice chat-bots, machine translation systems, but also underlying basic research).

Building an ecosystem of trust

Smart government approaches to regulation will play an important role in **boosting public trust and ensuring that AI is used responsibly, while also encouraging innovation**. However, it is important that a proportionate, risk-based approach is taken - balancing potential harms with the social and economic benefits that will be created by AI. The following recommendations aim to provide feedback on a number of important issues in the Commission's whitepaper.

Definition of AI

A clear and widely understood definition of AI will be critical to the effectiveness of any future regulatory framework. The White Paper describes the main elements that compose AI as data and algorithms. Such a broad framing effectively puts all contemporary software in scope; a narrower definition is needed to avoid over-regulation and to focus on the subcategory of AI systems where issues are most likely to arise.

Definition of high-risk

The Commission has rightly identified a **need for a well-defined risk-based approach to AI regulation** that doesn't apply 'one size fits all' logic across AI's myriad applications. With that in mind, a number of adjustments are needed to ensure that any potential regulation is targeted at the right use cases, provides legal certainty, and does not discourage the development and diffusion of AI:

- Factor in the opportunity cost of not using AI: It's vital that any risk assessment takes a holistic view, reflecting not only potential harms but also societal opportunities. The benefits will often outweigh the risks, especially if risks can be mitigated in a thoughtful way with strong safeguards. Regulation must not discourage use in such cases.
- Emphasise the need for proportionality: Risk assessments must reflect the probability of harm and not just the possible severity of the harm. It should also take account of the wider operational context when assessing risk, since the same AI application used for the same purpose will pose different risks

depending on the way it is integrated into business operations (e.g., extent of human oversight, additional safeguards such as monitoring).

- Remove open-ended statements that create legal uncertainty: The cumulative criteria (combining a list of sectors and clarity over what constitutes high risk use within them) is, broadly, a workable approach. However the "exceptional instances" clause that goes beyond the cumulative criteria (including AI applications that affect consumer rights) is too open-ended and creates legal uncertainty.
- Remove reference to "immaterial damages" in the risk definition: This is not a known legal concept, and should be clarified. As written it could mean anything from economic loss to hurt emotions, and could lead to legal uncertainty, discouraging investment and innovation.
- Be reasonable in the performance standards imposed on AI: There is a risk that innovative uses of AI could be precluded by demanding regulatory standards for AI systems far exceeding that required of non-AI approaches. While it is important to seek to minimise mistakes, no system, whether human or AI powered, will ever be perfect, and in some situations a lower level of accuracy may be acceptable. One example is a situation requiring an urgent and immediate response, where the cost of inaction is high and there are simply not enough qualified people on hand to do the job (e.g., helping triage medical screening in crisis settings). A sensible regulatory standard would be to require only that AI systems perform at least to a similar accuracy standard as would be expected of a qualified person carrying out a similar task, unless there is pre-agreed justification for an exception.

Mandatory requirements for high-risk AI applications

The White Paper outlines suggested mandatory legal requirements, a number of which could significantly hamper the development and diffusion of beneficial AI applications, especially if they are interpreted too literally. It is important to retain flexibility in the legal interpretation and to collaborate with practitioners to draft rules that are workable from a technical perspective. Key issues to address:

- Keeping of datasets should not be mandated:
  - Keeping datasets is likely to conflict with GDPR provisions requiring deletion of personal data, as well as presenting challenges for copyrighted datasets authorised for only short-term access.
  - It would destroy the privacy benefits of on-device processing because it would effectively force data to be collected and stored centrally.
  - It would prevent the use of off-the-shelf, open-source models, since developers will generally have no access to the data used to train them.
- Too much emphasis on training data quality; not enough on testing output:
  - The proposed obligations for developers to "ensure datasets are sufficiently representative" are impractical. They conflict with GDPR under which developers are not meant to have access to sensitive attributes like ethnicity. It is also unclear how to determine what is "sufficient" — especially for providers of multipurpose AI systems.
  - With enough expertise and care, it's possible to create a high-performing model even using biased, low-quality training data; the reverse is also true. Thus, rather than putting requirements on training data, it would be better to have requirements based on testing model performance using benchmark datasets, to make sure that the outputs are within an acceptable range, since it is the model output that ultimately determines the real world impact of an AI system.
- Do not take a literal approach to "reproducibility":

○ The whitepaper proposes "requirements ensuring that outcomes are reproducible". A too literal interpretation of reproducibility would be impossible to satisfy, as many AI systems have randomness built in, which makes it impossible to guarantee you get the identical output every time even if the input is the same. To be workable, there will need to be scope for broad notions of "predictability at scale" that do not require exact matching.

Enforcement

Ex ante conformity assessment requirements as recommended by the White Paper do not strike the right balance**. A combination of ex-ante risk self-assessment and ex-post enforcement for high risk AI applications would likely achieve similar results within much faster timeframes and without risking unduly stopping innovation and creating unnecessary burdens**. This would also build on existing industry practices, including the ethical, legal and due diligence practices that guide the responsible and trustworthy development of AI.

More generally, there are also effective and more workable alternatives to upfront conformity assessments that should be considered. Prior to any launch, for AI applications deemed to be high-risk, organisations could be mandated to carry out and document risk assessments based on articulated principles. This would be analogous to the requirement for data protection impact assessments under GDPR.

As currently proposed the ex-ante conformity assessment raises a number of issues:

● Treatment of products already in market: If it were deemed necessary for existing products in market to retroactively undergo conformity assessments, it would create a significant backlog for newly established testing centres. A grandfathering clause would solve this at the outset.
● Treatment of R&D and early-stage products: In the early stages of development there will often not be a clear view as to the ultimate shape of a product. It is therefore important that confidential testing and piloting of an AI application be allowed prior to any conformity assessment, within the bounds set by existing sectoral regulation. If such testing was not permitted, it may result in organisations taking an unduly precautionary stance in terms of the requirements and investment needed, which could hinder innovation. This would significantly weaken Europe's position vis a vis global competitors.
● Requirement to retrain on European datasets: The White Paper raises the possibility of requiring AI systems to be retrained using European data or in Europe, if developers are unable to prove that the original dataset used met European standards. This raises significant concerns. There is no guarantee that European datasets or training that takes place on European soil will do anything to improve the performance of an AI system. If anything, blocking the use of foundational (non European) data sets would risk reducing a system's performance, and could even exacerbate the risk of discrimination and lower quality outcomes. And such a requirement could cut Europe off from socially beneficial innovations developed elsewhere, for instance in the context of global AI research for COVID-19 solutions.

Safety and liability frameworks

Safety and liability frameworks must provide users of AI applications with sufficient protection, so if significant shortcomings are identified they must be addressed. However the Commission's report appears to conflate the notion of health and safety with concepts that fall outside the scope of product safety (e.g., cyber security, ethics, privacy and mental health). Any review of EU safety regulation should focus exclusively on areas where the unique properties of AI, IoT or robotics create a risk to the health and safety of consumers. To the largest extent possible this should be done at the level of special safety regulation (e.g. Regulation (EU) 2019/2144 on

type-approval requirements for motor vehicles). **There is no need for regulatory intervention related to general consumer devices where a security threat does not directly entail a safety issue.**

The current liability framework remains fit for purpose, being both effective and technology neutral, so sweeping changes are not needed. The long-standing provisions of European liability law have worked well, and there has been no showing of problems sufficient to warrant changing them and introducing the risk of unintended consequences. Any contemplated changes should be supported by clear evidence, and a strong consensus among legal experts, that the current framework is inadequate. Any changes should be consistent with the current framework and focus on consumer goods, but keep contractual arrangements as the norm in the B2B sphere and avoid any joint liability rules.

Globally, strict liability frameworks (such as the Product Liability Directive) are reserved for abnormally hazardous situations as they preclude any consideration of intent or negligence. The Commission should be wary of expanding the scope of the existing liability framework to cover AI software and services, as such a dramatic and unprecedented expansion would mean that anyone involved in making an AI system could be held liable for problems they had no awareness of or influence over. This would be particularly true if an operator decided to use the technology in a different way than intended by the developer (e.g. facial recognition might be used to guard public security but also for mass surveillance).

Another critical answer is the benchmark for comparison, which should be not perfection but instead improvements on current levels of human performance. There is a risk that faced with a sympathetic plaintiff, any judge or jury may assume that the creator of a new tool should be liable for anything bad that happens - not just performance at less than human level. We believe that if AI can improve over human decisions in many areas, but requiring that algorithms never make a mistake would deter any development of new algorithms.

Burdening AI system developers with such legal exposure would have a significant chilling effect on innovation and competition, one that would most likely disproportionately fall on European SMEs.

Signatories

Bulgaria

Croatia

Czechia

Hungary

Latvia

Lithuania

Poland

Romania

Slovenia

Slovakia