BUSINESSEUROPE

POSITION PAPER

15/07/20

AI: A European Approach to Excellence and Trust

KEY MESSAGES

- The aim of **achieving trust** in AI through policy needs to be **achieved with the goal of enabling its excellence** in Europe in mind.
- Europe needs to support its bid for AI excellence through **leveraging its digital frontrunners** to share best practices, coordinate policy actions holistically and spend resources in key areas that would support a sound basis for businesses to foster AI development and roll-out (eg. 5G, cybersecurity, data infrastructure, R&I, digital skills and standardisation).
- Achieving trust in AI through legal means should consider that it is a suite of technologies in its early stages. **Assessing existing laws and potential legal gaps** would be a good first step in order to adjust existing laws before new ones are made.
- The scope of any new requirements should take a risk-based approach and **only set market access requirements for "high-risk" AI**. This should be defined to focus on where the highest and most widespread societal damage is likely to arise. This is a pragmatic start to achieve trust while enabling AI development to continue to flourish.
- Legal certainty, specific responsibilities for all actors involved and a clear framework for business compliance in the delivery of AI need to be ensured so that AI or a product using an AI is **only covered by a single set of clearly assigned product safety rules**. As a result, either this new legislation for "high-risk" AI or existing sector specific legislation under the New Legislative Framework should apply.
- A **voluntary labelling system** for AI not covered under this new legislation **could be useful** to enhance trust. But each scheme should be defined following a bottom-up approach, identifying minimum criteria to be used by organisations choosing to participate in the same ecosystem.
- The potential legal gap of **new economic actors existing within "high risk" AI supply chains** that cannot legally be defined as "producers" in the context of the Product Liability Directive (PLD) should be explored.

AI: A EUROPEAN APPROACH TO EXCELLENCE AND TRUST

CONTEXT

BusinessEurope welcomes the positive tone of the Commission's White Paper acknowledging the many opportunities that AI can bring to Europe's economy and society. We support the Commission in its venture to build ecosystems of excellence and trust in Europe for Artificial Intelligence (AI). As an all-encompassing technology, this involves cross sectoral coordination across all areas of Europe through a number of legislative and non-legislative actions. No Member State alone can support our global competitive standing in this strategic technological area. Rather than relying on a few Member States or cities to lead, the diversity of resources which Europe can bring forward jointly in this process is what will propel it to be a global competitive player in AI. In return, the full economic, environmental, workplace and societal opportunities AI derives can be evenly spread across the continent - leaving no one behind through the ongoing digital revolution.

We welcome the Commission's pragmatic approach to excellence and trust in AI outlined in its recent White Paper. The two are likely to reinforce one another: a range of measures tackling legal challenges could support uptake by improving societal trust and offer greater clarity for consumers and businesses alike.

As a key political cornerstone of the Von der Leyen Commission, which will be fundamental to shape how AI develops in the EU, the utmost importance must be given to its transparent and open consultation before further operational steps by policy makers are taken. We repeat our call for an extension until the end of 2020 on this consultation in this regard due to the ongoing COVID-19 crisis. This will allow the fundamental debate required with stakeholders to take place. Adjusting timelines to the realities we are living within would ensure that quality prevails over speed. We suggest extending the period for input to the end of the year to ensure inclusive engagement so that useful proposals evolve from this consultation.

As a key societal stakeholder and European social partner (see latest digitalisation agreement), BusinessEurope outlines its reaction below to the Commission's White Paper on AI as part of this ongoing consultation.

1. AN ECOSYSTEM OF EXCELLENCE

The creation of an ecosystem of excellence begins with Europe's digital frontrunners who should seek to help other Member States move forward through sharing best practices and building an inclusive agenda that ensures digitalisation across the EU. The EU's Co-ordinated Plan on AI should be updated as a result to reaffirm commitment and coherence in this approach. Any legislation stemming from it should remain technologically robust in this rapidly advancing area.

Member States need to make greater funding commitments to develop innovative AI ecosystems. But the investment promised should above all be delivered in the EU's next long-term budget to offer the support needed at European level. AI investment in Europe is not sufficient. The €2.5 billion within the Digital Europe Programme was a first of its kind in this regard. The current delay and revision of the MFF is an opportunity to realise how important European development and uptake of new technologies such as AI is in battling the current COVID-19 health crisis (or any other) and how it can continue to aid Europe's long-term economic recovery if funded properly. While this was rightly recognised in the recent EU Recovery Plan Member States also need to move forwards to fully support this initiative also.

Europe needs top-class cyber-secure digital infrastructure to develop and run AI upon in order to foster our full capacities in this area. That means rapid and broad deployment of 5G that creates opportunities for all. Not just Europe's digital leaders or its largest cities. Member States should continue to implement the 5G toolbox, specifically enabling legislation related to risk assessments of suppliers and service providers. At the same time, we should begin to plan for harnessing 6G, to ensure Europe is better prepared for the next wave of digital infrastructure.

As AI is integrated into a broad array of systems, it will expand available entry possibilities for malicious actors to exploit. Therefore, it is important to support cybersecurity awareness' raising initiatives by promoting industry driven standards, guidelines and best practices to help companies manage their cybersecurity risks. The ENISA ad-hoc working group on AI should work with industry and come up with practical recommendations on how to mitigate AI risks and ensure that systems entering the single market are trained to respond to unforeseen circumstances.

For Europe to succeed in AI it must similarly succeed in the data economy. Incentivising greater voluntary data sharing and access, along with the opportunity for businesses to choose to use solid and reliable cloud infrastructure in Europe, will enable us to benefit from the data we generate, capitalise on the exponential growth of the global data economy and power our excellence in AI (a separate BusinessEurope paper details this further). The availability of trustworthy data infrastructure, based on the principles of portability and interoperability, is an important prerequisite for promoting a vibrant data economy.

More intense collaboration between government, regulators, enterprise and the research community to develop Europe's AI research and innovation ecosystem is needed. Supporting further European public-private partnerships in this area is crucial to consolidate our efforts. This collaboration could be supported further by public funding to support the proposal to create a lighthouse centre for AI research to attract and coordinate how AI can be applied in different sectoral areas. This approach must spread benefits across various geographical locations in the EU and not be centralised to one area. This could be achieved in a cost effective way if the existing network of Digital Innovation Hubs were leveraged to do so and expanded further, as proposed in the Digital Europe Programme. Permitting confidential testing and piloting of AI in the development stage should be permitted in any future framework to support Europe's AI research community.

Empowering people to understand and advance AI is key. Businesses should play a closer role in influencing all levels of national education so that foreseen labour market needs, such as embracing AI, can be linked closer to national curricula so that our citizens (including our workforce) gain the relevant STEM and transversal skills required to take part in the digital economy. Co-operation between the education sector and private sectors should be established and/or strengthened to support an education system that responds to future labour market needs. The reinforcement of the recent Commission skills agenda and the planned support of a network of education institutions offering world-leading masters programmes in AI are welcome. Citizens should also have a deeper understanding of how AI decisions are being taken. The EU needs to step up its efforts on raising awareness of the benefits which AI can bring in this regard. They could also help explain, with industry, how it can be used to help augment human involvement and capacities at work and actually improve allocation of tasks between workers and machines. Industry also has a role to play to continue updating the knowledge of its workforce as technology advances in this area, including those in wider ecosystems that will play a crucial role to support the uptake of AI. Some businesses have internal academies that provide programs for workers. Member States should encourage these activities further. We have found that there are three core roles (with corresponding skill sets) that are required within these programmes to make them a success:

- Developers (people who can create AI systems eg. AI experts as compared to domain experts);

- Trainers (people who can train AI systems eg. engineers who prepare and test data sets);

- Operators (eg. people who can operate AI systems and domains using them).

Europe's standardisation framework has a crucial role to play in fostering excellence in AI. Market-relevant technical standards can support interoperability, technology transfer and enable competitive levers to lead in AI applications. We continue to support an international approach in which European industry has an active and strong role lead the way. Europe should seek to lead where there are gaps in international standards and establish a "first mover advantage" (eg. as it did in GSM standardisation). European businesses need to better coordinate priorities of standards that support AI and determine when more of a European influence is required which can then be pushed towards influencing international fora. The Commission should think of ways to promote this coordination without interfering with the standard setting process itself.

2. AN ECOSYSTEM OF TRUST

Europe should incentivise trust in entire AI value chains without interfering with the efficiency of AI decision making itself. Enabling trust in AI through any new provisions should put transparency at the core. This is the main cross cutting societal issue that should be solved through the Commission's planned action. This should take into account:

- Consumer transparency: so citizens understand when an AI is being used, which functions are AI enabled, if any human oversight validation exists and where the responsibility for decision making could be placed;

- Business transparency: to trigger a positive feedback loop so that industry has transparency of the AI decision making process with as much clarity as appropriate. Clear responsibilities of all actors involved in the delivery of AI should be defined.

THE EXISTING AI FRAMEWORK AND GAPS TO BE CONSIDERED

As a collection of technologies and not a policy area itself, we note that AI is already regulated in several policy areas at European (eg. the General Data Protection Regulation (GDPR), the Product Liability Directive (PLD), the General Product Safety Directive (GPSD) etc.) and national level (eg. discrimination/equality legislation). These existing frameworks should be used first in a manner that promotes innovation whilst taking society with it. Only hereafter should we determine what legal gaps exist on the basis of demonstrable evidence to bring new provisions forward.

We accept that further legislation could be required in areas where current law as described above doesn't sufficiently answer societal questions that may become ever more apparent as AI develops. As a result, we believe there is a need to determine on the basis of evidence, whether updates are required in relation to some existing frameworks:

- Bias: the principles of the General Data Protection Regulation (GDPR) put individuals control as a priority, as a result, the data available for AI to learn from could be limited - reducing their ability to de-bias;

- Safety: as AI decision making can be more opaque in some cases than non-AI technologies;

- Liability: as the development of AI has included new actors in the supply-chain that cannot be defined as "producers" under the PLD.

A FUTURE REGULATORY FRAMEWORK

Any regulatory framework should take into account that AI as a collection of technologies is still in its early stages. This implies that while traditional policy instruments would provide legal certainty, they could also stifle innovation or slow down AI adoption. As AI increasingly comes to be seen as a necessity, the EU must consider the risks of *not* using it – for example, a medical application whose analysis could help speed up disease diagnosis. Therefore, while it is vital to minimise mistakes as far as possible, the EU must factor the cost of not using AI into policy debates.

To ensure legal certainty and ease business compliance with this new framework, a product or application should only covered by one set of rules on AI. Only "high-risk" applications should therefore be covered by the new horizontal regulation. However, if an AI application is linked to a product covered by NLF legislation, then the application should only be prone to the relevant sector specific regulation - not the horizontal AI rules. This with a view to ensure that products covered by more than one NLF legislation only have to comply with one set of requirements. New NLF regulation coming into force, which addresses AI (eg. software with incorporated AI) directly, such as the Medical Device Regulation, should be left untouched. Potential changes to the Medical Device Regulation should await future evaluations.

Further to this, as many aspects of AI in these existing sector specific areas are already covered (eg. software (with embedded AI) as a medical device under the incoming Medical Devices Regulation). Any new horizontal high-risk AI market access requirements should not inadvertently apply to those sector specific areas. At the same time, the high-risk criteria determined below should be used if policy makers determine that sector specific NLF legislation needs to be updated in the future.

We aim to find a level of risk that is acceptable to society. No existing technology or future innovation can demonstrate 100% safety. There will always be a residual element of risk. We should recognise that provisions should instead find optimum safety levels to enable Europe to benefit from the use of AI and enter global markets as leaders. Below we demonstrate a path to achieving this in relation to the Commission's current White Paper concept.

Definition of AI:

"Artificial Intelligence": should be defined as clearly and precisely as possible, as it will underpin the assessment of existing legislation, the design of future legislation, and their enforcement. Europe will be a pioneer in defining AI through regulation which could grant it competitive international influence. But we believe the White Paper definition proposed by the Commission is too broad (eg. describing AI's main elements simply as "data" and "algorithms" would include all contemporary software). This is a disproportionate and legally uncertain approach.

We believe it is vital to support a definition that insists on the human origin of any AI and highlights that a machine can only perform an action assigned from the outset by a human in any phase. We therefore support an adapted definition provided by the Commission's High-Level Expert Group on AI (AI HLEG):

"Artificial intelligence (AI) systems are software (and possibly also embedded in hardware) systems designed by humans that, given a complex goal, are taught by their designers or learn from experience how to act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions. As a scientific discipline, AI includes several approaches and techniques, such as machine learning (of which deep learning and reinforcement learning are specific examples), machine reasoning (which includes planning, scheduling, knowledge representation and reasoning, search, and optimization), and robotics (which includes control, perception, sensors and actuators, as well as the integration of all other techniques into cyber-physical systems)."

This definition can be backed up with additional information included in the report of the AI HLEG that supports it.[1]

[1] A definition of AI: Main capabilities and scientific disciplines High-Level Expert Group on Artificial Intelligence, AI HLEG (2019)

Scope of legislation:

Europe has already experienced uptake of more “narrow AI” that has been deployed effectively and safely for several decades. This has not caused a societal or legal need for existing legal frameworks to adapt and principle based, technology neutral frameworks have been robust enough.

More novel AI techniques that should be covered are those we can foresee coming to market in the next 5 years with a high societal impact. Therefore the scope of any new requirements should take a risk-based approach and only set market access requirements for “high-risk” AI.

This should be defined to focus on a category of AI systems where the highest and most widespread societal damage is likely to arise. This will determine which AI applications are covered in this regulatory step by the EU. This is a pragmatic start to achieve trust and excellence while enabling early AI development to flourish.

The Commission usefully attempts to define what “high-risk” AI is. We agree that it should take into account a number of steps to be determined (with some adaptation) to ensure provisions are targeted, legal clarity is supported and AI development is encouraged. This could be determined following the Commission’s current cumulative approach (with some alternations/additions) - the AI should only be deemed “high-risk” if both are affirmative:

(1) Is the AI in a sector that is deemed high-risk (to be defined in cooperation with the High-level Group on AI to remain robust)?

(2) Does the intended use of AI in this sector involve significant “material” risks (physical or damage to property only) that are likely to arise?

It is important to define AI applications as of high risk and not the entire sectors. We should keep in mind that any risk-based approach must be easily adopted by SMEs (unlike the GDPR’s accountability principle). Businesses will have to ultimately determine whether the AI they’re bringing to market, integrating, operating or developing is covered. As a result, specific support should be provided to businesses, particularly SMEs in order to reduce potential burdens that determination could create in practice. Further to this we would like to add that covering “exceptional instances” as defined by the Commission White Paper negates the entire benefit and clarity of the test they originally propose. Deeming any AI high-risk if it is used in the employment or consumer context would disproportionately overregulate this technology and bring a high-level of legal uncertainty to businesses attempting to determine whether they are covered.

Market Access Requirements:

Any market access conditions (requirements) set for AI should recognise that this technology is evolving rapidly and should adapt to global markets. In particular, it should be taken into account that this is a new technology that we will learn more about AI in the coming years. There should be enough flexibility as a result to adapt to new knowledge and evidence. At the same time, solving societal questions or utilising new innovation as fast as possible is crucial to ensure more companies will invest and experiment with new technologies within EU.

Training data:

The quality of AI output depends largely on the AI system, data that the AI system is trained on, and adjustments to the AI system based on testing that is carried out. But this doesn't mean that AI trained on "European" data will always bring about certain or even legally compliant results. While the stage of training AI with data is important, we should focus more on whether it creates legally compliant results on the basis of the data and its quality. With regard to the quality of data, reference to existing standards such as ISO/IEC 25024 "Measurement of data quality" and ISO/IEC 25012 "Data quality model" could prove useful.

Non-European data could bring different results when compared to European data, due to behaviours, habits, cultures, ethical approaches or conditions. This doesn't mean that it will be non-compliant with Member State or EU law (eg. discrimination law or the GDPR). Indeed, these societal nuances can be expressed also across the single market itself.

Instead, more efforts should focus on the output and operation of the AI and whether that in practice is compliant with EU law. In this way, Europe would not have to close its door to the use of non-European data to power its AI. Minimising the potential for businesses in Europe to choose different data sets around the globe or force them to "retrain" AI systems on European data could be self-defeating and risk greater discrimination. At the same time, to increase adoption of AI throughout value chains and support business strategies, manufacturers integrating or using AI need to know as much information as possible on how an AI system was trained. Encouraging AI developers to include this information to their business customers (eg. where from, how it was developed), without needing to disseminate any IP related information (eg. trade secrets), would support transparency and as a result, greater uptake of AI in Europe.

Keeping of records and data (conformity assessment):

We agree that keeping of records, documentation and where relevant, data sets need to be retained for a limited time period to ensure effective enforcement of this legislation. This would include carrying out a conformity assessment procedure of an AI application, which is linked to a product covered by NLF legislation, should be aligned with the NLF legislation. Moreover, the horizontal AI regulation for high risk applications should apply existing NLF modules in accordance with the level of risk.

But documentation of these risk assessments in relation to data to support conformity of the AI with relevant provisions should not have to be repeated simply because the product uses AI. This would greatly alter existing product safety frameworks in a confusing and disproportionate manner for all, particularly for SMEs. AI systems with machine-learning capabilities do not change functionality unless the manufacturer allows the developer of the AI system to make functional changes. Developers of AI systems define their scope within which they can "learn" and make "decisions" (e.g. optimize the functionality of the product). Keeping in line with NLF principles, it should only be up to the manufacturer to assess whether a substantial modification has been made and as a result requires a new risk assessment. This would be similar to the requirements of “traditional” machines, which under the machinery directive, requires new risk assessments if products are modified substantially.

Information provision:

In relation to setting further transparency requirements that go beyond record keeping we believe that citizens should understand when an AI is being used, which functions are AI enabled, if any human oversight validation exists and where the responsibility for decision making could be placed. However, demonstrating the “expected level of accuracy” may not be the most meaningful metric and therefore should be used with care. “Accuracy” will be specific to: the situation in which the AI is used (which will vary as similar products will be used in different environments), the various hardware in which similar AI is placed or differ depending on the data it is fed. Due to these variables, it is difficult for businesses to make blanket statements (eg. the level of accuracy of this AI powered product is 90%).

Robustness and accuracy:

Trust for AI will naturally flow from demonstration of its robustness and accuracy in the market. This could take place through developing systems with ex-ante consideration of the risks they may generate. However, determining the entire potential of the AI’s life cycle ex-ante doesn’t reflect AI which learns continuously. Many AI’s placed on the market learn from the end user. The influence and transparency shifts from the business considering risks at the ex-ante phase to the end-user or even operator. Therefore, pre-market requirements should not be placed on businesses to determine phases of the life-cycle that they have minimal impact over.

In some cases it would not even be desirable to determine the entire life-cycle before market entry. For example, autonomous vehicles are expected to use a high level of automation with self-developing AI. However, safety critical features will not evolve over time. The software responsible for critical safety in autonomous cars needs to be tested, analysed and correspond to standards to demonstrate functional safety before deployment. As a result, when the software is deployed, it is frozen and no changes can be made unless the manufacturer would update it.

The white paper also 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.

Bias:

Bias in AI decision-making may occur due to narrow or insufficiently diverse training data (eg. under-representing minorities), limited volume of training data (eg. due to opt-in/out approaches of the GDPR) or flawed algorithm designs.

While the potential for bias in AI systems is a growing societal concern, we believe that the GDPR and laws against discrimination are robust enough to handle these issues alone rather than create a new Commission initiative. We therefore support the European Data Protection Board's (EDBP) view in this regard. However, adapting specific GDPR Guidelines for application could go further to aid these concerns.

The development of AI is highly dependent on the availability of vast amounts of qualitative data. While algorithms are interchangeable, it is crucial to possess appropriate data to train the algorithm and subsequently create high quality AI-powered products and services.

The GDPR should champion opt-out or opt-in measures for various existing processing grounds (eg. legitimate interest or consent). But existing EDPB Guidance negates the use of contractual necessity as a ground to improve a service. As a result, the supply of personal data that reflects the true nature of society could be minimised due to the GDPR's support for the control of the individual. This could unknowingly lead to unbalanced datasets that economic operators have no option of de-biasing. Particularly when special categories of data (crucially needed to fight against bias itself) can only be processed on the basis of the explicit user consent under Article 9 of the GDPR.

Another obstacle that needs to be adequately assessed is the use of special categories of data. The processing of such data is essential to identify bias and at the same time is bound to strict limitations. Not being able to use such data to "de-bias" AI systems may result in a risk for individuals.

Human oversight:

Human input is central to an AI system's development. From problem articulation to data collection, data curation, product design, testing and monitoring, people are the engine in the creation of an AI. The objective of a trustworthy, human-centric AI can only be achieved by ensuring appropriate level of human oversight. But determining at what stage this oversight is appropriate in the delivery of an AI will greatly depend on the circumstances and what it is trying to achieve, even within the same type of product. For example, in the case of autonomous vehicles, there are different levels of automation and therefore different needs for human oversight. Some require the driver to intervene when it encounters a scenario the vehicle is not able to navigate and the driver must be available to take over at any time. Fortunately, the whitepaper correctly realises that the degree and stage of human oversight may vary from one case to another. Particularly as different uses of human oversight that are common-sense in one setting could be insufficient in another.

Selecting the most prudent combination through determining the form and stage of human oversight will rely on a holistic assessment of how best to ensure that an acceptable output is made. This will depend on the intended use and effects that it could have on safety and performance of the task at hand. An optimum level of safety on the basis of understandable trade-offs should be reached that is acceptable for society. Chasing "zero-risk" circumstances is unrealistic and will not foster excellence in European AI.

Voluntary labelling schemes

A voluntary labelling system for AI applications that do not qualify as “high-risk” could enhance trust in AI and foster its uptake. However, in order to be effective, labelling schemes should be defined following a bottom-up approach, identifying minimum criteria to be used by public and private organisations participating in the same ecosystem. This could mediate the interests among different stakeholders. Otherwise, dominant players could end up defining a label for all actors, including SMEs, and force it upon their value chains.

These schemes should not subject non-”high-risk” AI to the same mandatory requirements that high-risk AI must comply with. A scheme that uses stringent legal requirements aimed at high-risk AI is likely to be inappropriate for lower-risk applications, become highly onerous, disincentivise participation and fail to provide clarity for users. Instead, there should be broad agreement before such schemes could be feasible or helpful. The Ethics Guidelines could be relevant input here. But any such label should cover a variety of issues that work towards fostering trust on the whole (eg. not applying separate labels for ethics, privacy or security). Given the pace of change, any scheme would have to be outcome focussed to work as intended. Existing self-regulatory approaches should also be taken into account.

Liability:

While existing national and European level provisions are largely sufficient to provide an adequate legal framework, any need for revision should be clearly demonstrated. We agree with the most recent review of the Product Liability Directive (2018) in that it is generally fit for purpose but recognise some of the potential issues raised by the Commission report on the safety and liability implications of AI, the Internet of Things and robotics (2020).

Therefore, while the Product Liability Directive (PLD) should continue to apply to all consumer products, in order to achieve the right balance between innovation and consumer protection and to create trust in the uptake of AI in Europe, it could be necessary to clarify who is liable in case of an incident and for what they are liable for. For instance, the Commission could explore the potential gap in relation to players not currently defined by the PLD through its upcoming AI initiative with regard to players that influence high-risk AI but that cannot be legally defined as “producers” under the existing PLD.

As we have already demonstrated, transparency is crucial for the future of European AI. For business it will enable more sectors to understand and positively utilise AI systems. For citizens it would build greater public confidence to enable broader uptake. This should also achieve credibility of the entire market through clarifying who is liable in case of an incident and for what they are liable for. This would achieve greater legal certainty, fairness and trust in value chains.

As a matter of principle, all parties involved along the value chain should be covered, according to their individual contribution. The aim should be to fill unacceptable liability gaps in such a way that no party is unfairly burdened. It could be appropriate to cover next to “producers” (manufacturers) other potential players, those could be:

- Operators (providers that use the manufacturer’s product in a service provision and therefore have an element of control over it);

- Service providers (3rd party providers of software that affect how the AI will function and therefore have a degree of control so far as their services are used in the manner they were intended);

- Data suppliers (public or private organisation that creates, collects, aggregates and transforms data from public or private sources).

New provisions in the Commission’s AI initiative in relation to liability for high-risk AI could therefore widen coverage to operators and service providers (as described above). This in no way should make them solely liable for any incident that arises. Demonstrating a line of causality between the businesses defect in the performance of the AI and the harm caused would still remain crucial. As a result, we should only give the possibility to seek either: operators, service providers, or data providers liable if causality can be demonstrated within the remit of the products intended use (just as the status quo for manufacturers). Otherwise, an unlevel playing field will grow where the manufacturer becomes the only possible actor to be found liable for a product on the market when an incident arises, even if out of their control.

A better understanding of the existing possibilities under the PLD is needed in order to provide clarity on the concept of "product" – this could happen for example by continuation of the existing expert group on the PLD. In this context, we remind policy makers of the long-awaited guidance on the PLD that the expert group has provided input on. More clarity on the coverage of embedded software could also be enhanced through clearer guidance, including an update to the Commission Blue Guide in the area of product safety.

We should only give the possibility to seek either: operators, service providers, or data providers liable if causality can be demonstrated within the remit of the products intended use (just as the status quo for manufacturers). Otherwise, an unlevel playing field will grow where the manufacturer becomes the only possible actor to be found liable for a product on the market when an incident arises, even if out of their control.

In the accompanying report, the Commission suggests that reversing the burden of proof might be one solution to address the complexity of AI-related liability. It is vital that actors are liable only if causality can be demonstrated within the remit of the product’s intended use. Otherwise, this could be costly for SMEs and start-ups (a core component of the AI ecosystem) who both lack the capacity to prove that they had no responsibility for any harm and are least able to afford compensation costs.

Insurers must be closely consulted by the EU, to confirm that insurance remains available and affordable. This should however in no way become mandatory as “high-risk” AI is a host of technologies rather than a finite product (such as cars, for which insurance is mandatory).

We believe that any reference to “immaterial damages” should be removed from the Commission’s concept and forthcoming initiative on AI. Particularly as it is not defined in the White paper and a broad area with widely unknown legal concepts. Further to this, attempting to predict innovations that may enter mass markets in the next 5 years will create legal uncertainty, un-robust and unsuitable legislation. This would discourage investment and innovation in Europe. Only causing material damage (death, injury and property loss) should continue to be covered under this adapted liability framework.

CONCLUSION

BusinessEurope’s response to this consultation is only a first step in the support which we aim to grant all policy makers throughout this process in order to enable European excellence and trust in AI. We stand ready to discuss the actual details of any potential draft proposal or specific update to an existing legislation.

Without the correct policy and legal framework, the full potential of AI and as a result, the societal goals which it could achieve, will not be reached. This represents a global competitive challenge. As other regions race forward with large markets and relatively relaxed conditions to experiment within, their technological leaps will be greater and faster as a result. Europe faces a choice: get the balance between innovation and societal protection correct or accept and welcome a future where we will not have a choice in using European AI as the ecosystem of excellence to support it will not exist. Our aim to maintain trust in the entire value chain through policy needs to be achieved with the goal of enabling excellence in mind.