

European Commission Consultation on

**White Paper on Artificial
Intelligence:
A European approach to
excellence and trust**

Scottish Government Response

August 2020



Scottish Government
Riaghaltas na h-Alba
gov.scot

**European Commission Consultation on
White Paper on Artificial Intelligence
A European approach to excellence and trust
COM(2020) 65**

Scottish Government Response

Introduction

The Scottish Government welcomes the publication of the European Commission's White Paper on Artificial Intelligence (AI), and the opportunity to comment.

*The European Union's Strategic Agenda for 2020-2024: Scotland's Perspective*¹, published in January 2020, sets out how Scotland sees the EU's priorities for the period ahead and explains why these are of vital importance to Scotland and how Scotland can contribute to their delivery.

We continue to seek to collaborate with the EU across a wide range of issues, including mutually beneficial research and trade and investment, reflecting our shared values and goals.

In recent years, we have made significant investments in data-driven innovation and AI, and we are currently working with people and organisations across our country to devise an AI strategy for Scotland.

We have strong ties with the EU in this area, through collaborative research, business partnerships and supply chains, and as trade partners. We recognise the importance of the policies outlined in this White Paper for Scotland - particularly for our businesses who provide AI-powered products and services to the EU. We also see an opportunity for these proposals to further strengthen our ties.

We hope that our response to the consultation will provide a useful perspective, and form the basis for continued dialogue with the EU in this area.

¹ The European Union's Strategic Agenda for 2020-2024: Scotland's Perspective
<https://www.gov.scot/publications/european-unions-strategic-agenda-2020-24-scotlands-perspective>

Our shared values

AI offers many opportunities – for how we live our lives, our economy, our public services and society – that we need to harness thoughtfully and sensitively. As mentioned in this White Paper, used for good, AI can contribute to tackling the **Climate Emergency**², and achieving the **UN Sustainable Development Goals** – in Scotland, in the EU and globally. In addition, the role of AI has been and continues to be critical in the global response to the **COVID-19** crisis.

The Scottish Government has therefore committed to developing an **AI strategy for Scotland**.

The strategy will be aligned with our **National Performance Framework**³, which sets out the kind of Scotland we want to see and aims to give equal importance to economic, environmental and social progress – opportunities for all, increased wellbeing, sustainable and inclusive economic growth, and reduced inequalities. We were proud to host the **Wellbeing Governments Group**⁴ in 2019 with the Commission and others, and will continue to take a leadership role.

Our AI strategy will be published in 2021 and outline what Scotland needs to do over the next five years in order to realise the potential of AI, and **create the right ecosystem for people in Scotland to benefit from the adoption of AI as a trusted, responsible and ethical tool**.

We therefore wholeheartedly agree with President von der Leyen's statement that **AI must serve people**⁵, and support the Commission in its mission to ensure that AI is trustworthy, respects fundamental rights and does not cause harm. We share the Commission's vision of an ecosystem of trust as outlined in the White Paper, and would welcome a dialogue with the Commission to share insights from our respective engagements with citizens and organisations on how to achieve this vision.

The need for a new regulatory framework

We agree that there is a pressing need for a common European AI regulatory framework to address the risk of fragmentation in the common market. Such fragmentation would be to the detriment of both citizens and businesses, inside and outside the EU. A common, predictable and legally certain regulatory framework is essential to driving innovation and adoption of AI.

² Scotland's AI for Good Challenge to tackle the climate emergency

<https://www.scottish-enterprise-mediacentre.com/news/tackling-climate-crisis-through-artificial-intelligence>

³ Scotland's National Performance Framework

<https://nationalperformance.gov.scot>

⁴ Wellbeing Governments Group

<https://www.gov.scot/groups/wellbeing-economy-governments-wego/>

⁵ Press remarks by President von der Leyen on the Commission's new strategy: Shaping Europe's Digital Future

https://ec.europa.eu/commission/presscorner/detail/en/speech_20_294

We also see such a framework as an opportunity for the EU to be a **global standard setter** in AI – ensuring innovation and ethics go hand in hand. We welcome the EU’s active engagement in global fora including the UN and OECD to promote AI that respects “fundamental rights, including human dignity, pluralism, inclusion, non-discrimination and protection of privacy and personal data” – European values which Scotland shares.

We commend the Commission’s principle that “the new regulatory framework for AI should be effective to achieve its objectives while not being excessively prescriptive so that it could create a disproportionate burden, especially for SMEs.” The Scottish Government’s own Better Regulation agenda⁶ aims to reduce unnecessary burdens on business by ensuring all regulation is **proportionate**, consistent, accountable, transparent, and targeted only where needed.

To achieve this goal, we concur that “there is a need to **examine whether current legislation is able to address the risks of AI** and can be effectively enforced, whether adaptations of the legislation are needed, or whether new legislation is needed”, in this order. We welcome the investigative work already carried out by the Commission on the safety and liability implications of AI, as published in the companion Report to this White Paper. We would encourage this work to be widened to a comprehensive review of all potentially relevant current legislation, to inform future regulatory mechanisms.

We also agree that AI regulation must be both **agile** (“given how fast AI is evolving, the regulatory framework must leave room to cater for further developments”) and **pragmatic** (“any changes should be limited to clearly identified problems for which feasible solutions exist”). To help achieve those objectives, we would recommend that the Commission consider the use of sector-specific **regulatory sandboxes**. These would need to be carefully designed, taking into account lessons learned from existing sandboxes, and adequately resourced – and would therefore benefit from collaboration between regulators in the EU. The evidence gathered from such controlled experimental approaches could be useful to assess the need for both horizontal and sectoral regulation of AI. Conversely, sandboxes could help regulators “enhance their AI skills in order to effectively and efficiently implement relevant rules”, as proposed in the White Paper’s section on Skills, and help businesses establish the processes required to develop safe and trustworthy AI.

More broadly, we share the Commission’s working assumption that “the regulatory framework would apply to products and services relying on AI”, that is, not to AI *per se*. Given the vast range of potential applications of AI, many of which “we can only begin to imagine”, it seems prudent that new legislation required to address identified regulatory gaps be introduced progressively and from the bottom up, to avoid unforeseen negative consequences or obsolescence. **Sectoral regulation**, carefully tailored to each AI use case, built on the foundation of existing practice and expertise and further developed through the use of sandboxes where appropriate, could form the pillars upon which **horizontal legislation** would rest. This would allow the latter to focus entirely on addressing cross-cutting goals (such as upholding fundamental rights) and themes (such as data and decision-making) where these are not adequately covered by existing legislation.

⁶ Scottish Government Better Regulation agenda
<https://www.gov.scot/policies/supporting-business/business-regulation>

Risk-based approach

We agree that “a regulatory framework should concentrate on how to **minimise the various risks of potential harm**”, and that the risks identified, “both material (safety and health of individuals, including loss of life, damage to property) and immaterial (loss of privacy, limitations to the right of freedom of expression, human dignity, discrimination for instance in access to employment)”, are appropriate and relevant.

We note that the Commission proposes to implement this risk-based approach using the following process and mechanisms:

- Every proposed AI application (product or service) is required to be risk-assessed, and the White Paper implies that this should take place before development.
- If the AI application is deemed “high-risk”, it will need to meet a set of mandatory legal requirements in order to mitigate risk.
- Once the development of a high-risk application is complete, its conformity with the mandatory requirements will be assessed by regulators prior to its placement on the market.
- In case the conformity assessment reveals that the application does not meet the mandatory requirements, the shortcomings identified will need to be remedied, and presumably, conformity re-assessed.
- Once the application passes the conformity assessment, it can be placed on the market.
- Any prior conformity assessment will be without prejudice to monitoring compliance and *ex post* enforcement by competent national authorities. This would apply to high-risk AI applications as well as other AI applications subject to legal requirements, potentially with particular attention paid to the former.

We agree that those elements (risk assessment, mandatory legal requirements, prior conformity assessment, and monitoring and *ex post* enforcement) are likely to be necessary for implementing an effective EU regulatory framework. We also believe that successful implementation of the framework will require a parallel **investment in capability-building** for all the actors concerned, within the EU and in third countries.

We note the Commission’s view that the requirements should be “applicable to all relevant economic operators providing AI-enabled products or services in the EU, regardless of whether they are established in the EU or not.” We therefore particularly welcome this consultation, and believe it is important that the Commission continues to **engage with a broad range of stakeholders both within and outside the EU**, to ensure that the framework is effective and proportionate throughout its proposed global range.

Finally, we agree that **remote biometric identification**, for instance using facial recognition technology, comes with unique risks. It has come under scrutiny in the UK, and has been banned in certain US cities. We therefore welcome the EU’s intentions to further consider how best to use AI with biometric identification, and encourage the Commission to engage widely on this topic, exploring both public and private sector uses.

In addition, although the development and use of AI for military purposes is outside the scope of the White Paper, we would suggest that the Commission considers the risk for dual use (military, or terrorism) of broader categories of AI applications, such as computer vision and autonomous driving/flying, although we recognise that it will be challenging to mitigate this risk.

Risk assessment

We agree that “the determination of what is a high-risk AI application should be **clear and easily understandable and applicable for all parties concerned.**”

To achieve this objective, a two-factor, cumulative criterion is proposed in the White Paper for defining “high-risk” applications (in essence, risky sector, plus risky intended use). While this criterion is clear and easily understandable, its simplicity might also create challenges when screening certain complex AI use cases. Using the White Paper’s example (page 17), if an appointment booking system in the healthcare system failed or was erroneous (for instance, by assigning inappropriate priority to a patient’s appointment), this could have life-threatening ramifications on patients who require life-saving treatment – and therefore may be justified for “high risk” intervention, even though the intended use appears “low risk” when considered in isolation. There is also the possibility that that dual/multi-use applications could first be assessed in a “low-risk” sector, and subsequently repurposed for use in a different, higher-risk context, possibly via a software update.

The White Paper recognises that the cumulative criterion might not capture all risks and that “there may also be exceptional instances where, due to the risks at stake, the use of AI applications for certain purposes is to be considered as high-risk as such – that is, irrespective of the sector concerned”. We agree that that the purposes listed in the White Paper as an illustration (recruitment and workers’ rights, and surveillance) warrant particular scrutiny, but a purpose-based list of exceptions still would not capture the type of complex, context-dependent risky application illustrated in the medical appointment booking system example above.

More generally, we believe it will be challenging to differentiate in a binary fashion between high-risk and low-risk applications using a universal, clear and easily understandable set of criteria. Whilst helpful guidelines such as those proposed may serve to guide developers and regulators, we believe that risk assessment is likely to require an approach that is both finer-grained and more holistic. This will require careful examination and a degree of judgement, and considering that its outcome would have a significant impact on how an AI product is designed and brought to market, **assessment might require a dialogue between developer and regulator.** This would help ensure that, for instance, the developer of a low-risk application does not needlessly implement unnecessary requirements because it mistakenly assessed it as high-risk. This would also be necessary to avoid a situation whereby a developer (intentionally or unintentionally) incorrectly assesses their high-risk application as being low-risk, which under the proposed process might only be detected after consumers have been harmed. Such a dialogue would require resources from both regulator and developer, and the impact on SMEs in particular would need to be considered.

Ultimately, **risk is a continuum**, and it might not be necessary to find an optimal binary criterion. We believe that the mandatory requirements listed in the White Paper for high-risk applications represent good design principles in general and desirable outcomes which are relevant to any AI application. The level of risk for a given application would chiefly determine how stringently the application would need to meet those requirements.

The Commission also proposes establishing a **voluntary labelling scheme** for “no-high risk” AI applications – that is, not falling under the cumulative criterion or the exceptions listed in the White Paper, and hence not subject to the associated mandatory requirements. Whilst there are benefits to such an approach, we see some risks of potential consumer confusion which would need care to avoid. For example, the differentiation within “no-high risk” AI applications between those that carry the label, and those that do not might suggest to consumers that the latter category of applications could not only be of lower quality, but also potentially “less safe” – if consumers mistakenly interpret those labels as being similar to a CE marking.

Requirements and conformity assessment

As mentioned above, we agree that the proposed types of requirements (on training data; data and record-keeping; information to be provided; robustness and accuracy; and human oversight) are important; furthermore, we believe they are relevant for all AI applications, not only those deemed to be high-risk.

We fully agree with the Commission’s view that “the conformity assessments for high-risk AI applications **should be part of the conformity assessment mechanisms that already exist for a large number of products being placed on the EU’s internal market**. Where no such existing mechanisms can be relied on, similar mechanisms may need to be established, drawing on best practice and possible input of stakeholders and European standards organisations. Any such new mechanism should be proportionate and non-discriminatory and use transparent and objective criteria in compliance with international obligations.”

We agree that “particular account should be taken of the possibility that certain AI systems evolve and learn from experience, which may require repeated assessments over the lifetime of the AI systems in question.” We note that this could make conformity assessments a significant burden, and we would recommend the Commission considers the **potential impact on SMEs** in particular and how it could be mitigated.

Training data and algorithms

We agree that training data is fundamental to how AI machine learning applications perform, and that “measures should therefore be taken to ensure that, where it comes to the data used to train AI systems, the EU’s values and rules are respected, specifically in relation to safety and existing legislative rules for the protection of fundamental rights”. We

also agree on the importance of the three key considerations listed by the White Paper: safety, bias and discrimination, and privacy.

We note, however, that assessing the quality, diversity, and general fitness for purpose of a training dataset is a complex task, with implications for the skills that will be required from regulators and developers. **SMEs are at a double disadvantage**, in terms of access to large, high-quality datasets (as opposed to large incumbents such as Google and Facebook), and access to the skills and resources required to critically assess them and take remedial measures (such as correcting for bias).

It is also intrinsically difficult to assess whether a training dataset is “good enough”, in terms of how its characteristics will translate into an AI application making decisions that are safe and protect fundamental rights. Clearly, a biased training dataset is unlikely to lead to good real-world performance, or adequate protection of those rights. But apparent “diversity” in the training data is no guarantee that the resulting (for instance) neural network makes non-discriminatory decisions, or is generally fit for purpose. It would also be difficult to ascertain that a dataset covers all potential “dangerous” scenarios, as those are typically uncovered after the event. As pointed out in the White Paper, there will be a need to also verify the “relevant programming and training methodologies, processes and techniques used to build, test and validate AI systems” – not only the training data in isolation. This will place further demands on the skills of regulators, particularly if they are required to inspect the details of algorithms and underlying mathematics.

The implementation of **remedies** for training data will also pose challenges. The Commission suggests that a possible remedy is “re-training the system in the EU in such a way as to ensure that all applicable requirements are met.” However, it is unlikely that training a system in the EU would be either necessary or sufficient to meet those requirements. Such a remedy also raises the issue of how developers would access appropriate EU datasets for their AI application. Would such data be made available through the Commission, or would it be necessary for firms to procure such datasets from third party companies or collect data independently? This might put certain businesses, such as SMEs, at a disadvantage.

Keeping of records and data

The **accountability** principle is an important element of the GDPR. Organisations are expected to show how they comply with the legislation, for example, by keeping records of their processing activities and making them available to regulators on request. We believe that this requirement on AI systems could therefore be seen as an extension of those provisions.

To be effective, specific prescriptions to enable audit and accountability will need to take into account the complexities of machine learning systems. For instance and in broad terms, in AI applications based on neural networks, the system’s decisions are based on the pattern and weights of connections between the neurons in the network, which in turn result from training the network on a typically large dataset. The decisions made by the system thus reflect the underlying training data in indirect and complex ways. As a result, requiring the

developer to retain training data and make it “available upon request, in particular for testing or inspection by competent authorities” might not always be effective (in terms of how much it will help audit/forensic analysis) or proportionate (particularly for “deep” neural network AI applications, which require particularly large training datasets). In addition, as pointed out in the White Paper, machine learning systems often learn continuously, making training data a moving target. Finally, requiring data to be made available to authorities might raise intellectual property and trade issues, related to those that would arise from requiring source code to be made available, as well as data protection issues. The White Paper acknowledges such issues and proposes that “where necessary, arrangements should be made to ensure that confidential information, such as trade secrets, is protected”, but this might be challenging to achieve in practice.

Other requirements

We fully support the EU’s recognition of the importance of adequate and proactive **information provision** about the use of AI systems. We would welcome a wider exploration of transparency and explainability and what minimum standards would be appropriate, particularly in the use cases of statutory decision-making about individuals by public authorities, and life-critical decision-making for instance in health care. This is an active area of multidisciplinary research, including the development of technological solutions to better explain the decisions of neural network algorithms, and effective engagement with customers⁷.

Regarding **robustness**, we agree that consideration should be given to “requirements ensuring that AI systems are resilient against both overt attacks and more subtle attempts to manipulate data or algorithms themselves, and that mitigating measures are taken in such cases.” Adversarial AI, particularly for computer vision applications, is currently an unsolved challenge with major safety implications, such as for autonomous vehicles. We believe it is urgent to start providing legal clarity on what level of robustness would be appropriate for such AI applications, while bearing in mind this is a rapidly evolving field.

More generally, it is important to recognise that meeting the requirements listed in the White Paper will require **trade-offs**, for instance between performance (robustness and accuracy) and explainability/transparency, or between different, mutually incompatible definitions of fairness to achieve non-discrimination. For developers to comply and make appropriate trade-offs, it is critical that they are provided with clear and practical definitions of key terms (such as AI, diversity, human oversight and autonomy) and helpful **guidelines**. We welcome the work of the AI High Level Expert Group work in this area and look forward to seeing its updated guidelines following industry feedback. As previously mentioned, we believe there is value in encouraging best practice and upskilling in the key areas identified as part of the requirements, whether or not an AI application is high-risk.

⁷ UK Information Commissioner’s Office / Alan Turing Institute research and guidance on explainability <https://ico.org.uk/about-the-ico/research-and-reports/project-explain-interim-report>
<https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/explaining-decisions-made-with-artificial-intelligence/>

As discussed above for training data, assessment of conformity with the requirements will generally be a complex task. This will require regulators to develop skills in new areas and collaborate with other bodies. We therefore welcome the Commission's consideration of a "network of national authorities, as well as sectorial networks and regulatory authorities, at national and EU level" to facilitate the implementation of the legal framework.

Liability

We commend the Commission's effort to clarify **legal responsibility and accountability** for AI applications and systems. We agree with the Commission's view that, in principle, "each obligation should be addressed to the actor(s) who is (are) best placed to address any potential risks". However, complex situations might arise such as when AI systems are integrated into a product made by a third party, and/or receive continuous software updates. For instance, if large consumer businesses were not confident that suppliers providing the AI subcomponents of their products would bear appropriate responsibility in case their algorithm's failure resulted in the end product harming consumers, they might decide to bring most AI work in-house to better control risk. This could reduce competition and innovation, and particularly affect SMEs. We would therefore suggest that the Commission further explore how such legal clarity and confidence can be achieved.

Impact on SMEs and non-EU businesses

We welcome the Commission's specific consideration of the impact of the proposed regulatory framework on **SMEs**, which, as discussed above, could arise at several stages, and the proposal that "in order to limit the burden on SMEs, some support structure might be envisaged including through the Digital Innovation Hubs."

Since the Commission proposes that these rules apply to everyone providing AI solutions in the EU market, we would also urge the Commission to consider how the framework and enabling measures would affect differentially EU and **non-EU businesses**, for instance in terms of access to testing centres and support from Digital Innovation Hubs.

We welcome the emphasis placed on developing AI and digital skills in SMEs within the EU SME Strategy. The Scottish Government has been participating in the **OECD's Digital for SMEs Global Initiative**⁸, which aims to promote knowledge sharing and learning on how to enable all SMEs to make the most of the digital shift. This initiative places specific emphasis on the diverse opportunities and needs of the large "missing middle" of SMEs and on their role for an effective, inclusive and sustainable digital transition, and includes looking at the role of AI for SMEs. We will continue to engage with the OECD on three pillars that underpin the initiative: analytical research, sharing SME experiences, and network and policy dialogue.

⁸ OECD Digital for SMEs Global Initiative
<https://www.oecd.org/going-digital/sme/>

An Ecosystem of Excellence

We welcome the Commission's plans to "build an ecosystem of excellence that can support the development and uptake of AI across the EU economy and public administration". Our developing AI strategy's work strands to achieve this mirror those listed in the White Paper.

Scotland is a world leader in AI **research and innovation**. The University of Edinburgh is ranked first in the UK for computer science and informatics research. Universities, companies and others in and around Edinburgh attract research income from within the UK at almost four times the national average. The city region is home to unicorn technology companies as well as 400 start-up companies and a National Robotics facility⁹ enabled through the Data Driven Innovation Initiative¹⁰. The Scottish Government funds the Data Lab, Scotland's Innovation Centre for data and AI¹¹ to build upon these strong foundations and help realise the economic benefits of AI across Scotland.

Scotland is a strong partner in research and innovation **collaborations with the EU**. We have consistently argued for close ongoing involvement in EU research and innovation programmes, including the current Horizon programme and successor programmes¹². By April 2020, Scottish organisations had won over €50 million under the Industrial Leadership Pillar in Horizon 2020 for Information and Communication Technologies, the strongest topic for Scotland under this pillar. We will take appropriate measures to ensure a continuing flow of ideas, researchers and students between Scotland and the EU. We have already committed to extending free tuition fees for EU students commencing a course in 2020-21 at a Scottish university or college.

"Developing the **skills** necessary to work in AI and upskilling the workforce" will be a priority both for the revised EU Coordinated Plan and for our strategy. We are committed to providing opportunities for all to power our economy, and therefore wholeheartedly agree that "particular efforts should be undertaken to increase the number of **women** trained and employed in this area." We would be keen to exchange ideas and lessons learned with the Commission and member states on the most effective ways to achieve those aims, as well as **drive the adoption of AI and measure its impact** including in **SMEs and the public sector**. The Scottish Government has developed new ways for public sector organisations to solve policy challenges in collaboration with innovative technology businesses, such as CivTech¹³. We are leveraging and refining those approaches in our response to the COVID-19 crisis.

⁹ Scotland's National Robotarium

<https://www.edinburgh-robotics.org/robotarium>

¹⁰ Edinburgh and South East Scotland City Region Deal – Data Driven Innovation Initiative

<https://ddi.ac.uk/>

¹¹ The Data Lab

<https://www.thedatalab.com/>

¹² Shaping Horizons: Scotland's recommendations on the strategic planning for Horizon Europe

<https://portal.scotlandeuropa.com/file/download?id=3067>

¹³ CivTech

<https://www.civtechalliance.org/civtech>

Conclusion: looking ahead

The Scottish Government continues to be a keen supporter of the EU. Although we recognise that Brexit will unfortunately change radically how Scotland engages with the EU, Brexit will not change the EU's fundamental importance to Scotland nor our commitment to it. We will remain an active and positive partner who sees international collaboration as an essential driver for innovation and a healthy open economy.

EU regulations are of major interest to Scottish businesses exporting their products and services to a key trade partner, and to Scottish consumers who will benefit from the AI products and services made in the EU.

We hope that our response to the consultation will provide a useful perspective, and form the basis for continued dialogue and successful collaboration with the EU in this area, in accordance with the shared values that underpin the development of our AI Strategy and the EU policies on AI.



Scottish Government
Riaghaltas na h-Alba
gov.scot

© Crown copyright 2020

OGL

This publication is licensed under the terms of the Open Government Licence v3.0 except where otherwise stated. To view this licence, visit nationalarchives.gov.uk/doc/open-government-licence/version/3 or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

Where we have identified any third party copyright information you will need to obtain permission from the copyright holders concerned.

This publication is available at www.gov.scot

Any enquiries regarding this publication should be sent to us at

The Scottish Government
St Andrew's House
Edinburgh
EH1 3DG

ISBN: 978-1-83960-986-2 (web only)

Published by The Scottish Government, August 2020

Produced for The Scottish Government by APS Group Scotland, 21 Tennant Street, Edinburgh EH6 5NA
PPDAS754446 (08/20)

W W W . g o v . s c o t