


# An Intelligent Trust?

An Analysis of the EU White Paper on AI Regulation



Seamus Allen

## Key Takeaways: The EU Commission's White Paper on AI

1. The European Commission plans to promote AI investment and uptake in the EU through a variety of funding programmes. It is hoped that €20 billion - including both public and private investment - will be invested annually in AI technology for the next decade.
2. New regulations are proposed for “high-risk” AI applications only. To be deemed “high-risk”, an AI application must be used in a high-risk sector (e.g. healthcare, transport, energy) and must also have a function or use that is likely to give rise to significant risks. A voluntary set of standards may be offered for AI applications that are not deemed “high-risk”.

While a risk-based approach to AI regulation is welcome, the specific risk-based approach proposed by the White Paper is arguably lacking in nuance. This means that many AI risks will not face new regulations while other AI applications may be overregulated.

3. The new regulations involve five requirements for “high-risk” AI applications:
  - (i) The training data of the AI application must be sufficiently broad
  - (ii) Detailed records must be kept of the design, training and datasets used
  - (iii) Adequate information must be provided to potential users of the application
  - (iv) The application must be technically robust
  - (v) An appropriate level of human oversight is required

The five requirements are generally logical and constructive proposals. Some of these requirements should arguably be applied to AI technology in general, regardless of risk-level.

4. Prior conformity assessments will be used to test “high-risk” AI applications for compliance before they are placed on the market.
5. The regulations will be addressed to the actors best placed to deal with the risks posed by the AI application at a particular point in time. For instance, this may be the business operating an AI system rather than the original designer of the AI system.
6. The White Paper does not propose a framework for judging AI applications used in high-risk sectors where detrimental consequences for humans may be inevitable - for example, AI applications performing emergency decisions in healthcare or in autonomous cars. Without such a framework, high-risk sectors may be slow to adopt AI technology that may nonetheless be safer than and superior to its human counterparts.



# 1. Introduction

The European Commission's White Paper *On Artificial Intelligence - A European approach to excellence and trust*, which was published on 19 February 2020, is a bold step forward in EU strategic thinking on digital policy. This White Paper is open for public consultation and the closing date for responses has been extended to 14 June 2020<sup>1</sup>. Although the White Paper is not the Artificial Intelligence (AI) regulation which Commission President Ursula von der Leyen promised within the Commission's first one-hundred days, the White Paper nonetheless outlines a regulatory framework and strategy for AI and a plan to promote the development and governance of AI technology in Europe. The technical definition of AI adopted is: "Artificial intelligence (AI) refers to systems that display intelligent behaviour by analysing their environments and taking actions - with some degree of autonomy - to achieve specific goals."<sup>2</sup>

Underpinning this White Paper is a strategic desire by the von der Leyen Commission to strengthen the position of the EU in the geopolitics of the digital economy, which is currently dominated by the US and China, and to enhance Europe's digital sovereignty. The White Paper follows from the Commission's AI "strategy" and "Coordinated Plan", which were both adopted in 2018 to support the development and use of AI across Europe. The "Coordinated Plan" is due to be revised once the results of the public consultation on the White Paper are processed.<sup>3</sup> In order to ensure citizens' trust in AI, a High-Level Expert Group was established, which then published *Ethical Guidelines for Trustworthy AI* in April 2019. Revised guidelines are due to be released by the High-Level Expert Group (of which Prof. Barry O Sullivan, UCC, is the Vice Chair) in June 2020, which will take account of feedback received to date.<sup>5</sup>

## Mixed Responses to the White Paper

The regulatory proposals of the White Paper have been met with mixed reactions. Some technology law experts have hailed it as a "thoughtful and balanced way forward."<sup>6</sup> However, other responses have been much more critical; yet, the critics disagree with one another as much as they do with the advocates of the White Paper. *The Financial Times* has warned that "Brussels risks killing the technologies it wants to foster" as the proposed regulatory regime risks being too burdensome for business.<sup>7</sup> Eline Chivot, senior policy analyst at the Center for Data Innovation warns that the proposed approach risks to "deter or delay investment, the introduction and use of innovative products and services on the EU market[...]and lead to higher compliance costs for businesses launching AI systems in the EU".<sup>8</sup> However, Nicolas Kayser-Bril of AlgorithmWatch warns that the White Paper under-regulates and "fails to seriously appreciate the risks" of AI technology.<sup>9</sup> Corinne Cath-Speth of the Alan Turing Institute and Frederike Kaltheuner of the Mozilla Foundation similarly warn that the White Paper "falls short" by providing insufficient regulation and by having "missed a number of crucial opportunities."<sup>10</sup>

Such a wide disparity in opinion amongst the expert community is unsurprising for two reasons. First, some of the concepts in the White Paper are in need of greater clarification, as they have been interpreted in diverse ways by commentators. Second, the level of disagreement is also indicative of the unique regulatory challenges posed by the autonomous and dynamic nature of AI applications. Unlike regular goods or services, AI systems can change in character depending on the data to which they are exposed and do so with a certain level of autonomy that is independent of humans. As a result, AI technology both offers enormous potential to transform society for the better, and also poses unique regulatory challenges.



## 2. Supporting a European AI Sector

The White Paper first discusses the Commission's ambitious plans to promote the development and uptake of AI in the EU and to exploit the opportunities presented by AI technology. The Commission hopes that over €20 billion in total, in both public and private investment, will be invested in AI per year over the next decade.<sup>11</sup> It plans to draw on resources from a variety of programmes – among them the Digital Europe Programme, Horizon Europe, the European Structural and Investment Funds (ESIF) and the InvestEU programme – to promote AI investment and the uptake of AI technology in the economy.<sup>12</sup>

The Commission also plans to promote more networks and co-ordination in AI research and investment – between various organisations, economic sectors, the public and private sectors and third-level educational institutions.<sup>13</sup> It plans to facilitate the creation of “excellence and testing centres” where this co-ordination can take place,<sup>14</sup> in anticipation of the world's leading AI professors and scientists being attracted to Europe as a result, enabling Europe to provide world-leading master's programmes in AI.<sup>15</sup>

From a business perspective, the proposed promotion of skills development in the general workforce to facilitate the role of AI in the economy is welcome – for instance, through initiatives such as the Skills Agenda and the Digital Education Action Plan.<sup>16</sup> Of particular interest to the private sector is the proposal to ensure that AI will be accessible for SME's, particularly by supporting Digital Innovation Hubs and ensuring at least one hub is highly AI-specialised in each Member State.<sup>17</sup> The public sector will also be supported, with an “Adopt AI programme” to promote public sector AI procurement, which will particularly prioritise sectors such as healthcare.<sup>18</sup> Finally, the White Paper argues that AI regulation is itself another key pillar in the strategy for promoting AI technology in Europe, in order to ensure the trust of consumers, businesses and public services in AI technology.

## 3. The Risk-Based Regulatory Approach

The White Paper adopts a risk-based approach to the regulation of AI technology and divides AI applications into two categories – those AI applications that are deemed “high-risk” and those that are not. According to the White Paper, the proposed new regulatory requirements “would apply to high-risk AI applications only, thus ensuring that any regulatory intervention is focused and proportionate.”<sup>19</sup> It sets out two conditions to be fulfilled in order for an AI application to be deemed high-risk, both of which must apply in order for an AI device to fall into this category.<sup>20</sup> First, the AI application must be employed in a sector where significant risks can be expected to occur. Healthcare, energy and transport are provided as examples. Second, the AI application must be used in such a manner that significant risks are likely to arise.

### Conceptualising Risk

The White Paper understands risk as meaning the potential to produce “legal or similarly significant effects for the rights of an individual or a company; that pose risk of injury, death or significant material or immaterial damage.”<sup>21</sup> This seems to be a sensible and intuitive way of understanding the term “risk”. However, it adds that there may be “exceptional instances” in which certain AI applications can be considered high-risk even if they do not necessarily meet the above criteria.<sup>22</sup> The four examples specified are (i) the use of AI in recruitment processes, (ii) situations affecting workers' rights, (iii) remote biometric identification and (iv) other intrusive surveillance technologies; these examples will always be regarded as high-risk.<sup>23</sup> The Commission proposes to treat remote biometric identification (such as the controversial facial recognition technologies) separately from other AI technologies however, and plans to hold a “broad European debate” on this issue.<sup>24</sup>

In general, however, AI applications that are not deemed high-risk in accordance with the above criteria will not be subjected to the new regulations. It is suggested that they could participate in a voluntary scheme in which they adhere to a set of rules created specifically for the scheme (which are not discussed in the White Paper), or else to the regulations for high-risk AI.<sup>25</sup> In return they would receive a quality label to demonstrate their trustworthiness.<sup>26</sup>

In principle, such a risk-based approach to AI is a positive development, given concerns about how regulation could stifle the pace of innovation and AI-uptake. It seems reasonable that AI applications posing lesser risks should be subjected to less intensive regulation. Furthermore, it is welcome that AI applications used in high-risk sectors will not automatically be considered high-risk applications – the White Paper takes into account the functions of the applications involved. For instance, this means that AI-equipped entertainment devices in hospitals can be considered low-risk.

## Unregulated AI Risks

While the proposed “high-risk” regulatory approach in the White Paper is constructive, it is noteworthy that - as with any approach - there may be regulatory gaps. For instance, the White Paper aims for the proportionate regulation of AI in accordance with risk, yet only two categories of AI are proposed: a **high-risk category**, which will be fully and completely covered by the new regulations; and a **non-high-risk category**, which will be totally exempt. This is a rather imprecise and indiscriminate approach. AI applications in each category may vary widely in terms of the risks they may pose, and yet will be equally subjected to, or exempt from, the new regulations. This problem particularly arises for applications which pose a medium level of risk. These applications may either be unnecessarily over-regulated, if they are included in the high-risk category, or else may not be subjected to any new AI regulations at all.

Furthermore, it does not consider the contexts in which AI may operate as part of its risk assessment. Indeed, it may be far too complex to do so. However, this means that context specific risks may not be taken into account. While AI applications operating in a relatively risk-free sector such as cosmetics, or with a function that is generally low-risk (packaging, delivery logistics) may not be considered high-risk in accordance with the criteria above, such applications could still cause potentially major harm (for instance if a customer with a rare skin allergy receives an incorrect product). The White Paper itself states that an AI appointment scheduling system would not generally be considered high-risk, even if it were used in a high-risk sector, as it does not have a high-risk use.<sup>27</sup> However, the White Paper does not consider how, in specific contexts, such an application could still cause harm. For instance, if problems emerged with an AI scheduling system used for critical patients requiring check-ups, this could potentially have serious consequences.

Finally, the understanding of “risk” in the White Paper is centred on the possible detrimental impact to which an individual or legal entity might be exposed. As a result, Cath-Speth and Kaltheuner write that “many AI applications with far-reaching societal consequences fall outside the scope of the regulatory proposal.”<sup>28</sup> For instance, AI technology is now being used by social media companies to regulate online content (with many mistakes having already occurred).<sup>29</sup> AI is also being used in targeted advertising in ways that could create online echo chambers in which citizens only ever read content that supports their own political biases.<sup>30</sup> Both uses could have major consequences for public discourse, and could increase the polarisation of society in ways that are harmful for the functioning of a healthy democracy.



## 4. Key Problems with AI Identified by the White Paper

The White Paper denotes four key risks of AI technology – relating to discrimination, possible detrimental consequences, and regulatory and liability issues. These issues are the focus of the proposed regulatory and compliance strategies. The possibility of accidents or detrimental consequences, including those that may cause physical harm, being caused by AI systems is an important concern.<sup>31</sup> The White Paper also notes some prominent examples in which AI systems have been found to be responsible for discrimination, for instance, on grounds of race or gender.<sup>32</sup> Detrimental consequences could emerge from the use of AI applications for a variety of reasons. These problems could occur due to the design flaws of an AI application. However they could also arise due to poor datasets or flawed machine learning experiences, which are issues that could arise during the design phase but could also arise at any stage during the lifecycle of the application.

Crucially, therefore, AI applications can change over time and become problematic in ways that are not the fault of the original AI developer. From a regulatory perspective, it is thus difficult for authorities to know whether laws were broken during the design of an AI application, or whether detrimental consequences are the result of a flawed AI design process.<sup>33</sup> More generally, this can also give rise to liability issues. In most EU states fault-based claims need to be traced back to a specific person in order to be successful. This will be much more complicated with AI applications.<sup>34</sup>

## 5. The Five Requirements for High-Risk AI

The White Paper proposes five regulatory requirements which should apply to high-risk AI technology. These relate to (i) the training data used, (ii) the keeping of records, (iii) the information provided to

users, (iv) the robustness of applications, and (v) the level of human involvement. With regard to training data,<sup>35</sup> the White Paper proposes that data used to train AI products should be sufficiently broad while maintaining high privacy standards. Datasets should cover all relevant scenarios needed to avoid dangerous situations and to prevent prohibited discrimination (for instance, by ensuring datasets used are representative regarding factors such as race or gender). On the keeping of records and data,<sup>36</sup> the White Paper proposes that records should be kept to ensure problematic decisions by AI applications can be traced back and verified. This includes records and documentation relating to the training dataset and programming and training methodologies used. In the case of information provision,<sup>37</sup> it is proposed that “adequate information” should be proactively provided to citizens about the use of high-risk AI systems, including details on the accuracy, limitations, and conditions for effective functioning of the application. Citizens should be informed if they are interacting with an AI system rather than with a human. With regard to robustness and accuracy,<sup>38</sup> the White Paper proposes that AI systems should be technically robust, possess the designated level of accuracy and be developed with proper consideration of potential risks. AI applications must also be able to deal adequately with errors and inconsistencies and be resilient to both overt attacks and attempts to manipulate their dataset or algorithms.

On human oversight,<sup>39</sup> it is proposed that an appropriate involvement by human beings should be mandatory in relation to high-risk AI applications, although the appropriate type and degree of human oversight may vary from one case to another. This can range from cases where humans must make the actual decisions, to cases where humans simply have the ability to intervene, if necessary. For instance, AI applications can identify social security benefits applications that ought to be rejected, but only a human can make the decision to reject such an application. For AI-controlled autonomous cars, it is suggested that an emergency stop procedure which

a human passenger can activate would be adequate. The downside of this is that strict requirements on human involvement may reduce the potential benefits of AI systems to ensure safety, and clarification on these trade-offs needs to be provided.

While the five proposed requirements are welcome, it is notable that the White Paper does not propose a clear framework for judging high-risk AI applications. The proposed requirements themselves are a logical approach to AI regulation. Indeed, some of the requirements arguably ought to apply to all AI technology regardless of risk level – for instance the requirements that AI applications should be trained on sufficiently broad datasets and that AI designers should provide the adequate information specified. Conversely, requirements for resilience to cyberattack and data or algorithm manipulation are reasonable precautions for potentially life-threatening AI applications, but would be excessively burdensome for AI applications that are low-risk.

However, no framework or standard is proposed for judging AI applications involved in scenarios in which decisions with negative consequences for human beings may be inevitable (for instance, dangerous emergency healthcare procedures, or an emergency AI-car decision). There is also no framework to indicate if an AI application will be allowed to make “mistakes”. Without such a framework, AI developers may be deterred from creating AI applications that may be superior and safer than their human counterparts, as developers may fear being legally liable for the arising negative consequences. Such a framework is arguably essential for establishing an understanding of “trustworthy AI” in high-risk sectors.

## 6. Regulatory, Compliance and Governance Strategies

The White Paper proposes that the regulations will be implemented in a way that takes account of the dynamic nature of AI products, of the key risks identified, and of the EU’s desire to promote the

best possible AI regulatory standards and to project these standards and EU values globally. The methodology chosen is to use prior conformity assessments, to take into account the chain of users along which AI applications develop, to establish a specialised governance structure, and to maximise the geographic scope of the new regulations. The main compliance instrument proposed is the use of a prior conformity assessment – an assessment of high-risk AI applications before they are put on the market. This will include procedures for testing, inspection and certification, and verifying the programming, training data and the training and testing methodologies.<sup>40</sup> However, the dynamic nature of AI applications means that there is a limit to the possible effectiveness of static prior conformity assessments. The White Paper notes that “AI systems evolve and learn from experience, which may require repeated assessments over the life-time of the AI systems in question.”<sup>41</sup> Given this ability of AI applications to change over time, it is significant that the new regulations are to be addressed to the actors best placed to deal with potential risks at a particular moment in time. For instance, during the development phase this will be the designers of the AI application; at a later stage it may be the deployer, for example, a business that now uses this AI application to serve its customers.<sup>42</sup> However, the borders of responsibility may require greater clarification – for instance, on the degree to which AI developers must ensure that AI applications can filter inappropriate data versus the responsibility that later users of the application should have for ensuring that the application is not provided with inappropriate data.

With regard to geographic scope, another key proposal is that the new regulations should be imposed on all relevant economic operators providing AI-enabled products or services in the EU regardless of their origins. In other words, companies outside the EU wishing to sell AI-enabled products within the EU would have to be compliant with EU rules.<sup>43</sup> As is the case with the General Data Protection Regulation (GDPR), this requirement will have implications for third



country businesses, and EU standards and requirements may become a *de facto* standard for global business. To ensure adequate AI governance within the EU, it is proposed that a European governance structure would be established comprising the cooperating national competent authorities of the EU Member States.<sup>44</sup> This would facilitate the exchange of information, the identification of examples of best practice, and enable standardisation. It is proposed that maximum stakeholder participation should be encouraged – including from businesses, researchers, civil society, and consumer organisations.

## 7. Conclusion

Overall, many aspects of the approach in the White Paper can be considered positive and valuable steps towards building an AI regulatory framework in the EU. The principle of taking a risk-based approach and the key principles underlying each of the proposed five requirements – relating to training data, the keeping of records, information provision, robustness and human oversight – are constructive proposals. The decision that regulation will apply to the most relevant actor utilising an AI application, rather than just the developer of the application, takes into account the uniquely dynamic nature of AI applications. Prior conformity assessments to enable close inspection of compliance during the design stage of high-risk applications have met with general approval. Obligations of accountability, transparency and non-discrimination are dealt with at the level of principle.

The main flaw in the risk-based approach outlined in the White Paper is arguably that it is lacking in nuance. Many AI-risks will be left entirely uncovered by the proposed new regulations, even though some of the five requirements could be sensibly applied to them. For a variety of issues, greater clarification is needed about the approach which the European Commission intends to pursue.

The White Paper also lacks a framework of standards for assessing AI applications in scenarios where the deliberate decisions or unintentional mistakes of an AI application may give rise to negative consequences for humans, but in which the AI application could still be superior and safer than its human counterparts. Such a framework is arguably important to encourage AI-uptake in high-risk sectors. Ultimately, the Commission's White Paper represents a constructive first step towards AI regulation, but more thinking and consultation is needed in order to build a truly effective regulatory framework. As AI technology becomes increasingly important in both public and private sectors and in the daily lives of citizens, the regulatory framework emerging from the consultation on this White Paper will be of crucial importance for the future of Europe.

Geopolitical questions, however, remain to be answered. Although Europe is unable to compete at the level of "hard" high tech with the US and China, the question remains as to whether the EU will be able to become a global standard-setter for AI regulation, as it did with the GDPR in the area of data protection, thus projecting the "soft power" influence of the EU worldwide. In particular, it remains to be seen whether the post-Brexit UK will align itself with the EU's approach or will develop an alternative regulatory system of its own. The White Paper is nonetheless a key step in the development of AI technology in the EU and in the European Commission's ambition to make Europe fit for the digital age.





# Endnotes

- <sup>1</sup> <https://ec.europa.eu/digital-single-market/en/news/white-paper-artificial-intelligence-european-approach-excellence-and-trust>
- <sup>2</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 16
- <sup>3</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 4
- <sup>4</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 9
- <sup>5</sup> European Commission White Paper, “On Artificial Intelligence- A European Approach to excellence and trust”, p. 9
- <sup>6</sup> <https://www.lawfareblog.com/eus-white-paper-ai-thoughtful-and-balanced-way-forward>
- <sup>7</sup> <https://www.ft.com/content/7b894c24-59cc-11ea-abe5-8e03987b7b20>
- <sup>8</sup> <https://www.ft.com/content/7b894c24-59cc-11ea-abe5-8e03987b7b20>
- <sup>9</sup> <https://algorithmwatch.org/en/story/ai-white-paper/>
- <sup>10</sup> <https://tech.newstatesman.com/guest-opinion/eu-white-paper-on-artificial-intelligence-falls-short>
- <sup>11</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 5
- <sup>12</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 5
- <sup>13</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 5-6
- <sup>14</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 6
- <sup>15</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 7
- <sup>16</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 6
- <sup>17</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 7
- <sup>18</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 8
- <sup>19</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 18
- <sup>20</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 17
- <sup>21</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 17
- <sup>22</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 18
- <sup>23</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 18
- <sup>24</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 22
- <sup>25</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 24
- <sup>26</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 24
- <sup>27</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 17
- <sup>28</sup> <https://tech.newstatesman.com/guest-opinion/eu-white-paper-on-artificial-intelligence-falls-short>
- <sup>29</sup> <https://www.silicon.co.uk/e-innovation/artificial-intelligence/coronavirus-ai-takes-over-social-network-moderation-336399>
- <sup>30</sup> [https://www.huffingtonpost.in/ishtiaque-hossain/filter-bubbles-are-shrinking-our-minds\\_a\\_2146974/?guccounter=1](https://www.huffingtonpost.in/ishtiaque-hossain/filter-bubbles-are-shrinking-our-minds_a_2146974/?guccounter=1); <https://knightcenter.utexas.edu/books/NewInfoEnvironmentEnglishLink.pdf>
- <sup>31</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 12
- <sup>32</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 11
- <sup>33</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 13-14
- <sup>34</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 12-14
- <sup>35</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 19
- <sup>36</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 19-20
- <sup>37</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 20
- <sup>38</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 20-21
- <sup>39</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 21
- <sup>40</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 23
- <sup>41</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 23
- <sup>42</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 22
- <sup>43</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 22
- <sup>44</sup> European Commission White Paper “On Artificial Intelligence- A European Approach to excellence and trust”, p. 24,25



The Institute of International and European Affairs (IIEA) is Ireland's leading international affairs think tank. Founded in 1991, its mission is to foster and shape political, policy and public discourse in order to broaden awareness of international and European issues in Ireland and contribute to more informed strategic decisions by political, business and civil society leaders.

The IIEA is independent of government and all political parties and is a not-for profit organisation with charitable status. In January 2017, the Global Go To Think Tank Index ranked the IIEA as Ireland's top think tank.

© Institute of International and European Affairs, April 2020

Creative Commons License

This is a human-readable summary of (and not a substitute for) the license.

[https://creativecommons.org/licenses/Attribution-NonCommercial-ShareAlike 4.0 International \(CC BY-NC-SA 4.0\)](https://creativecommons.org/licenses/Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0))

You are free to:

- Share - copy and redistribute the material in any medium or format
- Adapt - remix, transform, and build upon the material
- The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

**Attribution** — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

**NonCommercial** — You may not use the material for commercial purposes.

**ShareAlike** — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

**No additional restrictions** — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.



The IIEA acknowledges the support of the Europe for Citizens Programme of the European Union.

