Response to the European Commission's Roadmap on Requirements for Artificial Intelligence

INTRODUCTION

On behalf of the Center for Data Innovation (datainnovation.org), we are pleased to submit feedback to the European Commission's roadmap titled "Inception Impact Assessment: Proposal for a legal act of the European Parliament and the Council laying down requirements for Artificial Intelligence."[1] In this submission, we summarize the four policy proposals under consideration and provide feedback on each option.

The Center for Data Innovation is the leading think tank studying the intersection of data, technology, and public policy. With staff in Washington, D.C. and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, cloud computing, and the Internet of Things. The Center is a non-profit, non-partisan research institute affiliated with the Information Technology and Innovation Foundation (ITIF).

BACKGROUND

The European Commission's white paper on AI, published in February 2020, argues that AI may cause harm and existing legislation on consumer protection, product safety, and liability does not adequately cover these risks. The Commission believes the EU should address these risks to ensure the development of lawful and trustworthy AI that respects fundamental human rights. To that end, the Commission is proposing four policy options, briefly summarized below:

Option 1: Soft Law Approach

This option would propose a flexible, non-legislative approach designed to facilitate, stimulate, and build on existing initiatives of industry stakeholders, such as AI principles, standards, and ethics

[1] European Commission, "Proposal for a legal act of the European Parliament and the Council laying down requirements for Artificial Intelligence" https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Requirements-for-Artificial-Intelligence.

guidelines. This option would also include monitoring and reporting on compliance with these initiatives.

Option 2: Legislation Establishing a Voluntary AI Quality Label

This option would introduce a voluntary labeling scheme that participants could use to demonstrate that their applications meet certain EU-wide requirements for trustworthy AI. These requirements would complement existing EU legislation that apply to AI systems, and would be based on the assessment list of the Commission's high-level expert group.

Option 3: Mandatory Requirements for All or Certain Types of AI

This option would use legislation to require AI systems to comply with criteria related to training data, accuracy, human control, and robustness, and would require record keeping for algorithms and datasets, and the provision of other information. These requirements would apply to a) specific categories of AI applications, such as remote biometric identification systems (e.g. facial recognition); b) all "high-risk" AI applications; or c) all AI applications.

Option 4: A Combination of Options 1, 2, and 3

This option would combine any of the prior options, taking into account the various levels of risk that could be generated by a particular AI application.

FEEDBACK ON PROPOSED OPTIONS

The Commission Has Failed to Make the Case for AI-Specific Regulations

As the Commission itself acknowledges in the roadmap, the potential harms posed by AI are not "new or otherwise necessarily tied to AI only." However, the Commission argues that action is warranted because EU regulators may be unable to effectively enforce rules on the protection of fundamental rights, product safety, and liability.

With respect to the protection of fundamental human rights, the Commission says that AI systems might make decisions with discriminatory outcomes which are difficult to perceive or challenge without documentation of how the system works. However, this same problem applies to discriminatory actions and decisions taken by humans which can similarly occur through opaque and

undocumented decision-making processes. Therefore, it is unclear why AI-specific rules are necessary.

With respect to product safety, the Commission notes that AI systems may not be covered by product safety legislation because some software is not explicitly covered by EU product safety legislation. But this is a problem that applies to all software-based systems, not just AI systems. Likewise, the argument that software updates may alter the performance of AI systems and thus require multiple conformance tests over the product life cycle applies to all software-based systems. Regulators should not seek to have one set of rules for AI-based systems and another for non-AI-based ones.

Finally, the Commission has not shown that liability questions about harms caused by software-based safety defects are any different for AI systems than for non-AI systems. If there are deficiencies in the EU's laws on product safety and liability, these should be addressed in a technology-neutral manner.

Since the Commission has not established the need for new rules targeting AI, it should not create them. Creating AI-specific regulations signals to consumers that the technology should be viewed with suspicion (thereby limiting demand), while the regulations themselves can make it harder and more expensive for companies to bring AI solutions to market (thereby limiting supply).

The hype around AI has taken on mythical proportions, but the Commission should adopt a clear-eyed view of this technology. AI has many important applications and its impact on the economy and society will be substantial, but it is not so fundamentally different from prior technologies that it necessitates an entirely new regulatory approach. Instead of trying to regulate AI systems, regulators should focus on oversight and accountability for how various organizations use AI systems.

The "Soft Law" Option Provides the Best Path Forward

Of the options outlined in the roadmap, Option 1—which would involve pursuing "soft law" options, such as promoting industry-led standards and codes of conduct, developing and sharing best practices, and providing guidance across sectors—is the best approach. This will create opportunities for businesses to develop AI solutions and address potential concerns, and for policymakers to identify specific gaps, if any, in existing EU legislation. The use of regulatory sandboxes for AI in specific sectors could support this process. It is too early to know how AI will develop, how organizations will use it, how consumers will respond, what problems might occur, and how different stakeholders will respond to those problems. It is better for regulators to not apply the precautionary

principle to AI, and instead to wait and see how the technology, business models, and consumer practices develop before moving forward with additional regulation.

A Labelling Scheme Would Be Premature and Could Have Unintended Consequences

Option 2, a voluntary labeling scheme, presents several risks.

First, it would create additional compliance costs and administrative burdens for businesses developing AI solutions making it more difficult for them to bring AI products and services to market. For example, organizations may have to pay a third party to certify that their AI solutions meet the requirements of a particular label. Larger businesses would likely be at an advantage over smaller ones since they have the resources to participate in a labeling scheme. Moreover, some businesses may choose to eschew using AI solutions in their products to avoid a label, hurting AI development in the EU.

Second, it is unlikely that industry stakeholders will agree on a one-size-fits-all labeling standard for AI given the diverse range of AI-enabled products and services across sectors. A standard that is too weak would be unhelpful while one that is too stringent would keep legitimate products off the market. Industry-specific labeling standards are more likely to emerge organically where they are useful.

Third, establishing a label for AI could suggest to consumers that unlabeled AI solutions are dangerous, thereby depressing adoption of these products and services. Creating an AI labeling scheme would likely increase prices for AI solutions, leading to greater economic divide between consumers. Moreover, AI-specific quality labels only address one part of quality and ignore the fact that, in some cases, the absence of AI may be an indicator of a subpar solution.

Finally, the long and generic assessment list of ethical principles which would form the basis for this labeling scheme overlooks the feedback of the private sector, which raised concerns over the requirements this list includes, such as explainability and transparency, their redundancy, and the difficulty companies will have to implement them in practice.[2] The Commission should therefore revise this list as these concerns were not address in its latest update.

[2] Eline Chivot, "Initial Lessons Learned From Piloting the EU's AI Ethics Assessment List" (Center for Data Innovation, March 1, 2020), https://www.datainnovation.org/2020/03/initial-lessons-learned-from-piloting-the-eus-ai-ethics-assessment-list/.

Mandatory AI Regulations Would Hurt EU Competitiveness

Option 3 to create mandatory requirements for AI would make it harder for the EU to compete in the digital economy. As explained previously, the Commission has not provided evidence of the need for AI-specific regulation. Historically, the EU has wisely chosen not to regulate software tools—like spreadsheets or computer chips—even though failures of these products could have significant consequences. Instead, it regulates particular sectors or activities, such as financial services or financial reporting. It should be no different with AI.

Existing EU laws and regulations, including sector-specific ones, already provide a framework appropriate for AI systems, and the Commission has yet to show evidence of any fundamental legal gaps that are unique to AI. For example, the EU Machinery Directive and the Product Liability Directive are technology-neutral and already formulate safety requirements that apply to machines using AI systems, including those that did not exist when these directives were adopted. Moreover, creating explicit references to AI in these directives would likely become outdated quickly as the technology evolves. Instead, policymakers should continue to refine sector-specific regulations to address novel uses of AI, such as updating regulations on vehicles to allow for the use of autonomous vehicles.

Option 3a would create specific rules around remote biometric identification systems such as facial recognition. Here again, there is no need to regulate the technology itself, but instead it would be more prudent to limit specific uses of the technology, such as preventing mass surveillance by the government or requiring law enforcement use of facial recognition to meet specific performance requirements. Existing laws, including the General Data Protection Regulation (GDPR), already apply to the use of these systems. Policymakers should be careful not to pursue policies that would prevent organizations from using the technology for legitimate and safe purposes, such as to increase convenience and security for consumers.

Option 3b, which would only establish new requirements for “high-risk” AI applications, could be a viable option if the Commission only considers it for high-risk scenarios where there is clear evidence of consumer harm, not just hypothetical harm, and these harms are unique to AI application. So far, that threshold has not yet been met. In addition, this option could require companies to seek prior approval from regulators before bringing AI systems to market. Unless the EU has sufficient testing capabilities, these reviews would prevent or delay the introduction of new AI systems. Moreover, high-risk AI applications are likely sector-specific and would be addressed by sector-specific regulations, such as for motor vehicles.

Option 3c, to establish EU-wide AI requirements for all AI systems, is the most problematic in that it would extend regulation to cover the whole spectrum of AI systems, from the simplest to the most advanced, especially as AI becomes ubiquitous. Such a broad scope would likely force some commercial AI applications off the market due to the costs involved, to the detriment of countless businesses and their consumers.

Finally, Option 3 considers requiring human oversight of AI systems. It will be difficult for regulators to assess the appropriate degree of human oversight of AI systems. Some systems may not need to have a "human in the loop" to perform safely and effectively, and creating a requirement for dedicated human review of algorithmic decisions would make some AI systems impractical and inefficient. Organizations should be given freedom to decide the degree of human oversight that is necessary based on the context in which they deploy an AI solution. One way policymakers can give businesses this flexibility goal is by creating a regulatory framework that penalizes organizations for causing harm based on the degree of harm and their intent. Smaller penalties should result when an organization does not harm consumers and acts unintentionally, while larger penalties should result when an organization harms consumers with its actions and that organization was either willfully negligent or acted with malicious intent. By promoting algorithmic accountability, rather than dictating how algorithms work, regulators can provide effective oversight and redress for consumers, while allowing businesses to make decisions about how to best manage risks presented by AI.

Finally, EU policymakers should strive to use AI to increase EU productivity, which has grown at anemic rates. To that end, they should pursue AI policy options that encourage automation and do not unnecessarily require redundant human oversight.