

Response of Microsoft Corporation to NIST RFI on an Artificial Intelligence Risk Management Framework

15 September 2021

Introduction

Microsoft Corporation (“Microsoft”) welcomes the opportunity to respond to NIST’s request for information on the development of an AI Risk Management Framework (“AI RMF”). As a developer and user of AI, Microsoft is dedicated to empowering every person and organization to unlock its vast potential. Key to this will be ensuring that AI is developed and deployed in ways that are responsible and trusted by users. As such, we are supportive of NIST’s effort to develop the AI RMF which can be an important part of building trust and advancing innovation in AI technologies. We look forward to contributing to the AI RMF’s development.

As NIST has recognized, the AI RMF should be developed through a broad, consensus-driven, open, and collaborative process that will include workshops and other opportunities to provide input. Microsoft shares NIST’s goal of developing a voluntary framework that consists of outcomes and processes that align policy, business, technological, and legal approaches to improve risk management.

A key challenge for NIST will be to create an AI RMF that is coherent with the existing NIST Cybersecurity and Privacy Frameworks. In addition, the AI RMF should be created in such a way that it is flexible and adaptable to advances in the technology, maturation of the practice of responsible AI, and the evolving landscape of policy and regulations associated with AI systems. To achieve that goal, Microsoft encourages NIST to pursue an approach to the AI RMF that is: (1) risk-based and outcome-focused, (2) cognizant of the broad constellation of AI technologies, associated use cases, and the need for nuanced analyses that rigorously assess benefits and costs, (3) responsive to the different roles of different actors across the AI value chain and the sociotechnical nature of AI technologies, (4) interoperable with existing similar frameworks and international standards, and (5) forward-looking. This approach will enable future implementers of the AI RMF to establish safeguards that secure the benefits of AI systems, manage their risks, and balance global AI policy regimes, regulatory obligations, and actionable engineering choices.

Response to the RFI

1. The greatest challenges in improving how AI actors manage AI-related risks—where “manage” means identify, assess, prioritize, respond to, or communicate those risks;

Reflecting the roles and responsibilities of different actors in the AI ecosystem

With its benefits, AI can also create risks. These sociotechnical risks concern the capabilities and limitations of the technology, combined with people’s expectations of it and the societal context of its use. For example, risks may emerge at the intersection of system design decisions taken by AI model and decisions taken by deploying organizations as to how and where and when to use the AI model in a final system. Many AI systems provide generally applicable functionality, like text analytics¹ or anomaly

¹ Example of Microsoft’s Text Analytics service: <https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/>

detection², meaning that a customer can decide to use them in a wide variety of scenarios, often in combination with AI offerings from different suppliers, as part of a larger system. As such, the deployer of an AI system, with their understanding of the specific use case, is the actor best placed to ultimately identify and mitigate risks which will be specific to their chosen scenario. The developer of a system, with their knowledge of its design and capability, should cooperate with the deployer so they can make informed deployment and risk mitigation decisions. Given the broad range of AI systems deployments and the varied nature of supplier and deployer dynamics, the appropriate allocation of responsibilities between actors will differ from use case to use case. Finding a way to appropriately reflect the roles and responsibilities between these actors will be a key challenge that the AI RMF will have to address.

Adopting a flexible approach focused on outcomes and supporting processes

Microsoft supports NIST's goal of ensuring the AI RMF is adaptable across the AI ecosystem and is grounded in an approach focused on outcomes rather than one-size-fits-all requirements. The AI RMF must be flexible enough to stretch to the breadth of AI use cases and remain effective over time in the face of rapid technological development and maturation in the practice of responsible AI. To achieve this, the AI RMF should attend to providing guidance on the outcomes that an organization should look to achieve in addressing AI risk, supported by a description of the approaches that can be taken to do so. As NIST intimates in the proposed attributes for the AI RMF, this approach would be preferable to setting out a series of fine-grained requirements, for example around specific dataset composition, that will likely prove ineffectual and quickly outmoded in the face of technological change and ongoing development of risk mitigation techniques and tooling. A balance will be required here between optimizing for flexibility and providing enough practical guidance to help organizations identify, measure, and mitigate risk. This will be particularly important for organizations that are in the earlier stages of building out their responsible AI processes.

Furthermore, it is critical to bear in mind that AI spans a broad constellation of technologies and that approaches to risk and outcome-focused analyses may differ depending on the technology employed and intended uses of AI offerings. Given the broad use of AI in complex domains of application, conceptions of risk and its characterization and management may need to extend to rich cost-benefit analyses, where suboptimal and costly outcomes are characterized and tolerated given the overall expected value delivered by a system.

Focusing on highest risk use cases

Microsoft supports NIST's goal of ensuring an AI RMF is risk-based and believes the AI RMF should focus on identifying and mitigating risks associated with the highest risk use cases, e.g. those adversely affecting legal rights or life opportunities, (such as access to healthcare or education); significant physical or psychological injury; and harms to fundamental rights. This will help focus mitigation efforts and resources on highest risk use cases, helping ensure they can be successful without impeding the use of AI technologies for other lower risk scenarios.

Adopting a lifecycle approach

AI systems are dynamic, with many continuing to adapt throughout their lifecycle as they learn from the data they process, as are the societal contexts into which they are deployed. Moreover, AI is developed, operated, and maintained as a service rather than as a fixed product. As such, it is important to adopt a lifecycle approach to monitoring and responding to the risks of a particular deployment given the way its performance may change over time. This includes developing processes for identification and mitigation

² Example of Microsoft's Anomaly Detector service: <https://azure.microsoft.com/en-us/services/cognitive-services/anomaly-detector/>

of risks during both the design and deployment of an AI system, including ongoing evaluation of the impact of a system throughout its deployment. Assessment of risk over the lifespan of an AI system must address changing behaviors, workloads, and associated outcomes due to both (1) updates to data, models, parameters, and overall functionality of deployed systems that may come via maintenance and updating, and (2) changes in the nature or distributions of tasks or workloads analyzed or handled by the system over time.

Supporting organizations in developing governance frameworks

Governance frameworks play an important role in helping organizations identify and mitigate AI risk. These frameworks should provide for a set of requirements, practices, and training to ensure that those developing and deploying AI are able to think through the impacts of these systems such that risks can be identified and mitigated over time. Governance frameworks should also ensure that management structures are developed to prioritize and respond to AI related risk.

Developing and scaling these processes within an organization is not without some complexity and thought should be given to how to help smaller organizations do this effectively. It is also important to note that practices around responsible AI are still nascent. Progress has been made in understanding and identifying risks, however more work is needed to develop measurement norms and effective mitigation techniques for the full range of AI's use cases and potential harms. For example, significant advancements have been made around how to identify, measure, and mitigate quality of service harms³ for computer vision systems. However, progress in other areas has been more limited. For example, many of the questions around how to identify and mitigate representational harms⁴ are still at the research phase and more work is needed to build out practical approaches that organizations can utilize.

Supporting compliance and engineering stakeholders

Alongside the development of governance frameworks, it is important to appropriately enable the internal communities that operationalize a framework. Two important communities are compliance teams and engineering stakeholders. Compliance teams are responsible for developing the practical internal controls that support and measure fulfilment of a governance framework's requirements. Engineering and associated stakeholders (e.g., software developers and data scientists) are responsible for ensuring that as they are building and implementing AI solutions, they are doing so in accordance with the framework. Thinking through the investments needed to effectively operationalize an organization's governance framework so that it can be easily implemented by these internal communities, and ultimately by the engineers building out solutions, will be important if a governance framework is to succeed.

- 2. How organizations currently define and manage characteristics of AI trustworthiness, and;**
- 3. How organizations currently define and manage principles of AI trustworthiness:**

Questions 2 and 3 are interrelated and our response to both are combined below:

Adopting a principled approach

Given the wide-ranging nature of AI use cases and associated risks, it is important for any approach to AI governance to be grounded in a set of guiding principles. At Microsoft, we have six AI principles that

³ Quality of service refers to whether an AI system works as well for one person as it does for another, even if no opportunities, resources, or information are extended or withheld

⁴ Representational harms occur when a system reinforces stereotypes or demeans or erases certain demographic groups

underpin our responsible AI program: fairness, reliability & safety, privacy & security, inclusiveness, transparency, and accountability⁵.

Organizations around the world have developed their own sets of principles in recent years, including the Organisation for Economic Co-operation and Development (OECD) which has developed a set of intergovernmental principles⁶ that we encourage NIST to consider in developing the AI RMF (more detail in [section 7](#)). Globally, a large number of principles have now been developed (Stanford's 2021 AI Index Report referenced 117 sets of principles created globally between 2015 and 2020⁷) with significant overlap between offerings and a growing consensus around the fundamental principles that should be prioritized. Resonating with conclusions and recommendations of the final report of the National Security Commission on AI, high-level principles need to be transformed into specific practices⁸. Thus, NIST's focus should therefore be on developing practical measures that can help organizations move from principles to practice in addressing AI risk.

Microsoft's responsible AI program

Microsoft continues to build out its responsible AI program, which is designed to ensure that the company is developing and deploying AI in ways that uphold our AI principles⁹. Our learnings may be informative to NIST as it develops the AI RMF. The program has a number of key elements, including:

- **Responsible AI governance:** Our responsible AI governance approach builds on, and is integrated with, the company's frameworks for privacy, security, and accessibility. We employ a hub-and-spoke model, with the hub comprising three groups: Aether, the Office of Responsible AI, and RAISE. Aether, which stands for AI, Ethics, and Effects in Engineering and Research, includes working groups that leverage top scientific and engineering talent to provide subject-matter expertise on the state-of-the-art and emerging trends regarding the enactment of Microsoft's responsible AI principles. Our Office of Responsible AI sets internal policies and governance processes and enables and coordinates the effort across the company; and the Responsible AI Strategy in Engineering (RAISE) group which works with engineering teams to utilize technical frameworks and tooling for responsible AI. The spokes of the model include our Responsible AI Champs who sit in engineering and sales teams across the company and raise awareness internally about responsible AI issues as well as help teams spot and address related issues.
- **Microsoft's Responsible AI Standard:** The standard is a set of rules to guide colleagues around how to enact Microsoft's AI principles in their work. The standard adopts an outcomes-focused approach setting out concrete goals that teams designing and deploying AI systems must adhere to, supported by a set of implementation processes. Teams conduct an impact assessment for each system to identify the purpose of the system, what its impacts might be, including on stakeholders who may be indirectly affected by the functioning of the system, and what the benefits and harms of the system may be. Teams are required to document this process to allow for review and traceability.

⁵ <https://www.microsoft.com/en-us/ai/our-approach>

⁶ <https://oecd.ai/ai-principles/>

⁷ https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report_Master.pdf

⁸ <https://www.nscai.gov/2021-final-report/>

⁹ <https://blogs.microsoft.com/on-the-issues/2021/01/19/microsoft-responsible-ai-program/>

- **Sensitive uses process:** While governance processes are important in addressing AI risk, in the fast-moving and nuanced practice of responsible AI, it is impossible to reduce all the complex sociotechnical considerations into an exhaustive set of pre-defined rules. This understanding led us to create a sensitive uses process for ongoing review and oversight of high-impact cases. The process requires that use cases meeting review criteria (which includes systems that may adversely affect legal rights or life opportunities—such as access to healthcare or education, result in significant physical or psychological injury, or pose a threat to human rights) undergo a review process to ensure they can be delivered in a way that adheres to Microsoft’s AI principles.

4. *The extent to which AI risks are incorporated into different organizations’ overarching enterprise risk management – including, but not limited to, the management of risks related to cybersecurity, privacy, and safety;*

Ensuring interoperability with existing frameworks

The AI RMF should be designed such that risks related to AI can be incorporated into existing risk management frameworks. This includes aligning to ISO 31000 Risk Management¹⁰ and ISO 23894 AI Risk Management¹¹ (more detail below in section 5). Building on and aligning to NIST’s Cybersecurity and Privacy Frameworks will also be important to ensure coherence across these different frameworks. In particular, there may be an opportunity to extend the concepts in Section 2 of the NIST Privacy Framework to the tasks of understanding and managing AI risks.

To allow for effective alignment with existing approaches, it will also be important to clarify the type of risk that the framework is intended to address i.e., is it organizational risk, or risk to users or individuals? ISO standards, for example, traditionally deal with organizational risk; however, impact to individuals can also be included in enterprise risk calculations, as is the case with privacy risk management frameworks.

Addressing the sociotechnical nature of AI risk

While aligning the AI RMF to existing standards and frameworks is an important goal, the framework will also have to be responsive to the sociotechnical nature of AI risk and the way in which this differs from privacy and cybersecurity risks. Unlike many other technologies, AI services often offer generally applicable functionality that is deployable in a wide range of different scenarios. Each scenario poses its own set of risks that are influenced significantly by the societal context into which the system is deployed. For example, a restaurant customer could choose to deploy Microsoft’s text analytics¹² service to analyze customer reviews for positive feedback, a relatively low risk deployment. The same restaurant could use the same service to scan CVs for key words as part of shortlisting job applicants, which could pose a greater risk of harm. As such, the risks that an AI system can pose are heavily shaped by the decisions that a customer takes in deploying the system, often beyond the visibility of the designer of the system, as well as the decisions taken by the developer around system design. The significance of the customer’s deployment decisions is even greater when a customer “fine-tunes” a customizable AI model by conducting secondary training of the model using its own data, to further localize it to its chosen scenario. For the AI RMF to be successful, it must be responsive to the respective

¹⁰ <https://www.iso.org/iso-31000-risk-management.html>

¹¹ <https://www.iso.org/standard/77304.html>

¹² <https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/>

roles and responsibilities for developers and deployers of AI and the ways in which AI risk differs to other risks, including cybersecurity and privacy.

5. Standards, frameworks, models, methodologies, tools, guidelines and best practices, and principles to identify, assess, prioritize, mitigate, or communicate AI risk and whether any currently meet the minimum attributes described in the RFI;

Relevant existing international standards and frameworks

We welcome NIST's stated aim of designing the AI RMF so that it is adaptable and consistent with other approaches to managing AI risk. An important part of this will be ensuring the AI RMF is informed by the significant amount of work that has been done in recent years to create common standards and frameworks for identifying and mitigating AI risk. This will include reflecting and driving alignment around existing international standards, including the many existing technology standards that are applicable to AI even if they were originally developed for other technology segments, including data formats, transfer protocols, cybersecurity practices, privacy practices and cloud services practices.

There are a number of existing standards that help meet attributes described in the RFI that NIST should be mindful of in developing the AI RMF. ISO/IEC 42001¹³ AI Management System in particular provides a strong foundation for a governance, risk and compliance framework in an organization. ISO/IEC 23894¹⁴ AI Risk Management, ISO 31000 Risk Management¹⁵, ISO/IEC 38507 Governance implications of the use of AI by organizations¹⁶ and the COSO (Committee of Sponsoring Organizations of the Treadway Commission) Enterprise Risk Management Integrated Framework¹⁷ are also important to consider in developing the AI RMF.

More broadly, ISO/IEC 22989 AI Concepts and Terminology¹⁸, ISO/IEC 24028 Overview of Trustworthiness AI¹⁹ and ISO/IEC 19944-1, Cloud Computing and Distributed Platforms – Data Flow, Data Categories and Data use – Part 1: Fundamentals²⁰ are important to note, given the way they provide standardized definitions of key terms and concepts.

Responsible AI documentation

Responsible AI documentation plays an important role in helping ensure that developers and deployers of AI systems can make informed choices about their design and use. To this end, Microsoft publishes Transparency Notes²¹ for several of its platform AI services, setting out the capabilities and limitations of each service, as well as considerations for responsible use, such as the impact that a deployment environment might have on performance. We see important synergies between our Transparency Notes and other industry efforts such as Model Cards²², Datasheets for Dataset²³, and AI FactSheets²⁴.

¹³ <https://www.iso.org/standard/81230.html>

¹⁴ <https://www.iso.org/standard/77304.html>

¹⁵ <https://www.iso.org/iso-31000-risk-management.html>

¹⁶ <https://www.iso.org/standard/56641.html>

¹⁷ <https://www.coso.org/Documents/COSO-ERM-Executive-Summary.pdf>

¹⁸ <https://www.iso.org/standard/74296.html>

¹⁹ <https://www.iso.org/standard/77608.html>

²⁰ <https://www.iso.org/standard/79573.html>

²¹ <https://docs.microsoft.com/en-us/azure/cognitive-services/responsible-use-of-ai-overview>

²² <https://modelcards.withgoogle.com/about>

²³ <https://www.microsoft.com/en-us/research/publication/datasheets-for-datasets/>

²⁴ <https://www.ibm.com/blogs/research/2020/07/aifactsheets/>

Microsoft is also participating in the Partnership on AI's ABOUT ML²⁵ initiative to evolve the artifacts and processes for responsible AI documentation across the key components and development lifecycle stages of machine learning systems, bringing together multiple approaches via a synthesis. We applaud the multiparty stakeholder approach to the development and refinement of the ABOUT ML documentation recommendations.

Impact assessments

Impact assessments have proven their value in a range of fields, including in relation to data protection, and can play a vital role in helping organizations address AI risk. Impact assessments can help those designing AI systems think through the impact these systems will have on individuals and society more broadly, and ultimately help identify effective mitigation methods for any identified risks. They also play an important role in documenting the steps that are taken to address any identified risk, helping drive accountability for responsible design and deployment, and are a critical element of any AI governance framework. Similar to the need for the AI RMF to be coherent with the Cybersecurity and Privacy Frameworks, we recommend that NIST's approach be informed by an understanding that organizations undertake management and engineering practices, including crafting impact assessments, in partnership with teams often performing impact assessments across a number of different risk areas. Thus, AI impact assessments related with the AI RMF should be coherent with impact assessments in other domains.

Encouraging the use of responsible AI tooling

Greater use of responsible AI tooling will be an important part of addressing AI risk. There are a growing number of tools that can help developers identify, diagnose, and mitigate risks that may be emerging from their models, including the following open-source tools that Microsoft has helped develop:

- Fairlearn²⁶: A toolkit that allows developers to assess model fairness and better identify tradeoffs in their model fairness and performance. It also includes a set of unfairness mitigation algorithms that help address the fairness issues identified in a model.
- InterpretML²⁷: This utilizes interpretability techniques to help developers better understand the behavior of their models with a view to identifying and mitigating any responsible AI issues.
- Error Analysis²⁸: This helps engineers analyze model errors across different subgroups within a dataset to understand where trust, reliability, and fairness issues may be arising.
- Counterfit²⁹: A tool that helps engineers proactively find security vulnerabilities in their AI models before they can be exploited by adversaries.

6. How current regulatory or regulatory reporting requirements (e.g., local, state, national, international) relate to the use of AI standards, frameworks, models, methodologies, tools, guidelines and best practices, and principles;

²⁵ <https://partnershiponai.org/workstream/about-ml/>

²⁶ <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>

²⁷ <https://interpret.ml/>

²⁸ <https://erroranalysis.ai/>

²⁹ <https://github.com/Azure/counterfit/wiki>

Aligning the AI RMF with developing regulatory frameworks

To ensure that the AI RMF is consistent with other approaches to managing AI risk, we encourage NIST to remain engaged with, and aligned to, the development of regulatory frameworks elsewhere. The ongoing development of the EU AI Act is of particular importance in this regard given the way in which it will create a comprehensive, horizontal regulatory framework for AI. Microsoft supports the goals of the EU AI Act proposal and the way in which it adopts a risk-based framework focused on regulating the highest risk use cases. Microsoft also supports the way in which the Act provides for self-assessment of conformity and adherence to harmonized standards, and the way it provides for voluntary codes of conduct. We also believe there are ways to strengthen the Act, including by ensuring regulatory obligations fall on the entity best placed to meet them and by adopting more of an outcomes-based approach. We set out our views in more detail in our formal response to the European Commission's AI Act proposal³⁰.

7. AI risk management standards, frameworks, models, methodologies, tools, guidelines and best practices, principles, and practices which NIST should consider to ensure that the AI RMF aligns with and supports other efforts;

Aligning with OECD principles

To help meet NIST's goals for the AI RMF, including around adaptability and driving consistency with other approaches to managing AI risk, we encourage NIST to align to the OECD's AI principles³¹, designed to ensure that AI can be innovative, trustworthy, and respectful of human rights and democratic values. The OECD principles offer a set of intergovernmental principles for trustworthy AI that have been adopted by OECD member countries, including the United States Government³², and provide a strong foundation for the AI RMF.

Aligning with existing standards and frameworks

As outlined above, Microsoft recommends aligning the AI RMF to existing standards and frameworks, including the NIST Cybersecurity and Privacy Frameworks. Alignment is also important with regard to other AI risk management standards to ensure an effective and efficient risk management regime.

The BSA recently published an AI risk management framework, entitled a *Framework to Build Trust in AI*³³, which may also be a useful resource for NIST to consider. The BSA framework articulates a lifecycle-based approach to addressing AI risk through the use of impact assessments and mitigation techniques which are important elements of any approach to managing AI risk.

³⁰ https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/F2665556_en

³¹ <https://oecd.ai/ai-principles>

³² <https://www.state.gov/artificial-intelligence/>

³³ <https://ai.bsa.org/confronting-bias-bsas-framework-to-build-trust-in-ai>

8. How organizations take into account benefits and issues related to inclusiveness in AI design, development, use and evaluation – and how AI design and development may be carried out in a way that reduces or manages the risk of potential negative impact on individuals, groups, and society.

Important components of responsible AI development

Microsoft believes the following are important elements of developing and designing AI systems in a way that secures their benefits and reduces the risk of potential negative impacts of the technology:

- **Creating diverse and multidisciplinary teams:** Creating diverse product design and development teams that represent a range of perspectives and disciplines is important to ensuring systems are developed in a way that is inclusive and reflective of the stakeholders that will be using them.
- **Conducting impact assessments:** A structured assessment of the impact a system will have on stakeholders is important in identifying and mitigating any sociotechnical issues that a system may create. This stakeholder assessment should assess impact not only on users of the system but other stakeholders across society, including any marginalized groups that may be differentially impacted by the system. Relevant stakeholders should also be consulted at the different stages of an AI system's lifecycle.
- **Setting outcomes-based goals:** Creating outcomes-focused goals in a governance framework can help advance responsible outcomes. This can be done advance fairness goals, for example, by requiring AI systems used in high-risk deployments to provide a similar quality of service and a similar allocation of resources or opportunities to groups impacted by the system. More detail can be found in section 2 of our formal response to the European Commission's AI Act proposal³⁴.
- **Prioritizing training and transparency:** It is important to provide adequate training and information to those designing and using AI systems so that they understand how to design and use them in a way that is responsible. As part of this, the developers of AI systems should make available information to customers about the system's capabilities and limitations, to help them make informed deployment decisions and identify and mitigate any risks in their chosen deployment scenario. As highlighted above, Microsoft currently provides this type of information to customers of our platform AI services via Transparency Notes³⁵.

³⁴ https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/F2665556_en

³⁵ <https://docs.microsoft.com/en-us/azure/cognitive-services/responsible-use-of-ai-overview>

9. The appropriateness of the attributes NIST has developed for the AI Risk Management Framework

The importance of a risk-based, adaptable, and interoperable framework

Microsoft believes that NIST has set out an appropriate set of framing attributes for the AI RMF. In particular, it will be important for the AI RMF to be developed in a way that is consensus-driven, adaptable to the breadth of the AI ecosystem, risk-based, outcome-focused, and voluntary. As outlined above, an important part of this will be supporting organizations to develop a consistent set of requirements, practices, and training so that those developing and deploying AI are able to do so in a responsible way.

We welcome that the attributes outlined in the RFI align to the approach used for both the NIST Cybersecurity Framework and the NIST Privacy Framework. Microsoft also appreciates the goal of ensuring the AI RMF is capable of being updated as technology and risk mitigation techniques develop. We would encourage NIST to ensure that any updates are made in the same open, inclusive, and collaborative manner that helped make the Cybersecurity and Privacy Frameworks a success. An open and inclusive approach will also help ensure that changes to the AI RMF are informed by an up-to-date understanding of the technology and the use cases to which it is being put, as well as the maturing practice of responsible AI and evolving sociotechnical risks.

10. Effective ways to structure the framework to achieve desired goals, including integrating AI risk management processes with organizational processes for developing products and services for better outcomes in terms of trustworthiness and management of AI risks

The AI RMF should build on existing NIST frameworks

As outlined above, interoperability with existing standards and frameworks will be important for the AI RMF to meet its goals. In particular, the AI RMF should align to the structure of NIST's Cybersecurity and Privacy Frameworks which have proven effective.

11. How the framework could be developed to advance recruitment, hiring, development, and retention of a knowledgeable and skilled workforce necessary to perform AI related functions within an organization

Standardized frameworks can help grow demand for AI roles

The existence of a robust AI RMF will help build trust in the technology, encouraging deployment of AI and growing demand for people skilled in the wide range of disciplines involved in creating and using AI responsibly. This includes those working in technical roles, as well as those with a background in humanities and social sciences who play a critical role in identifying and addressing the sociotechnical risks that AI can pose. Standardized frameworks for responsible development and use will also help with the important task of ensuring organizational leaders are equipped to oversee responsible AI governance programs as part of utilizing AI in their organizations. There is also the potential for NIST to advocate for the creation of more standardized, specialist roles in organizations in relation to AI governance and compliance.

12. The extent to which the Framework should include governance issues, including but not limited to make up of design and development teams, monitoring and evaluation, and grievance and redress.

Governance frameworks are an important part of mitigating risk

Microsoft believes that the AI RMF should address governance issues. As outlined above, governance frameworks are important in ensuring that appropriate organizational measures and accountability mechanisms are put in place, particularly where the practice of responsible AI is still maturing and norms as to the identification, measurement, and mitigation of risk are still developing. These frameworks may include the specification of roles and responsibilities for adherence to the principles, the creation of oversight processes for high-risk systems, and requirements to develop appropriate policies, operating procedures, and control systems to ensure that risk management practices are consistently implemented.

Conclusion

Microsoft welcomes the opportunity to share the above comments which are informed by our experience of developing and deploying AI with customers and building out our responsible AI program. We applaud the work that NIST is doing to build out the AI RMF and look forward to contributing to this process, alongside others. We look forward to opportunities to participate in upcoming forums or workshops on this and other topics related to the responsible use of AI.