



April 29, 2022

Indeed's Response to the National Institute of Standards and Technology's Request for Information on an *Artificial Intelligence Risk Management Framework*

Indeed's mission is to help people get jobs.

Since inception in 2004, Indeed has put the Job Seeker experience at the heart of every decision we make. With 75M unique visitors per month¹ in the US and 250M globally², more people find jobs on Indeed than anywhere else, allowing Job Seekers to search millions of jobs in more than 60 countries and 28 languages. Globally, over 3M employers use Indeed to find and hire new employees.

As our company has grown, this mission has never changed, but our means to accomplish this goal have evolved to include tools that utilize machine learning, algorithms, and other advanced technologies. Some of these tools could be characterized as “artificial intelligence (AI),” though it bears noting that there is no consensus on the meaning of that term. These tools open up new opportunities for us to better match job seekers and employers. We invest in product solutions and advanced technologies to help make the lives and jobs easier for those who use our platform. For a small or medium sized business who uses Indeed to find great people, connect with their top picks, and to make job offers, our technology investments, including AI and algorithms, free up their time so that they can spend their time and money to help their customers and improve their products. As well, we are committed to using our resources to drive innovations that reduce barriers many job seekers face during the hiring process.

Indeed's product solutions and advanced technologies are helping to facilitate our goals of making job search faster and simpler, removing barriers to employment, and establishing a standard for measuring and improving work happiness.

While we recognize the potential of these technology tools to further drive us to our mission of helping, we understand the potential risks they present and are fully committed to using AI in an ethical and responsible manner, and in the service of driving inclusivity. Our use of these technologies is done in accordance with awareness and understanding that there are limits to

¹ Comscore, Unique Visitors, February 2022

² Google Analytics, Unique Visitors, February 2020

the technologies and controls must be put in place to ensure the ethical development and use of such tools. We seek to structure our advanced technologies through the deployment and active use of guiding principles:

Inclusivity and Fairness - We are determined to be a fair and inclusive platform for employers and job seekers to connect.

Privacy - Preserving the privacy, security, and trust of job seekers and employers is our top priority.

Transparency - We are committed to providing a clear and accessible explanation of how our technology works.

Humanity - Humans and diverse perspectives are integral to the building, use and review of AI tools. Our tools help expedite and increase the accuracy of our users' decision making.

Learning from Best Practices - We are constantly evaluating our AI technology, determining how it can better serve our users.

We greatly appreciate and encourage the value and effort of NISTs work to establish a framework which can be deployed in the design, development, use and evaluation of AI products, services, and systems.

We recognize that the field of AI regulation and policy is increasingly crowded, with various jurisdictions, stakeholders, and practitioners offering AI guidance, frameworks, recommendations, and best practices. As such, any US framework or frameworks should be established in context of existing efforts underway, like the [OECDs AI principles](#).

That said, the very nature of AI and similar advanced technologies is their evolving nature. Any effort to structure, regulate, or provide frameworks to AI must be created with the objective of not merely representing a static snapshot of the current state of AI, but rather be structure in a manner which allows flexibilities and growth, to accommodate rapid and future developments and changes to the technology and its deployment for good.

In the context of risk and AI, the role of development of any framework should be centered around the principle of mitigating risk, and not the impossible objective of complete elimination of risk.

Indeed offers the below selected comments in direct response to NISTs request for input on specific areas related to its draft risk framework:

Coverage, Specificity, and Flexibility

The general framework represents a broad perspective, perhaps in an attempt to be applied to a wide variety of scenarios and applications. A general perspective presents some difficulties in the application or use of such principles. One-size-fits all solutions would offer little operational benefit except to provide a basis where future, more specific or segmented principles can be developed.

We support NIST's further development, in coordination with relevant stakeholders, of industry or AI application-specific principles.

Enablement of Organization Decision Making

Fundamentally, we appreciate the role the RMF plays in the framing of risk as a potential adverse impact as well as the likelihood that it will occur. In broader debate and discussion around AI, a tone of absoluteness can occur, which in the practical application of AI tools is untenable.

The RMF could benefit from further explanation of the tradeoffs that may occur between mitigating risk and operational effectiveness of AI tools, as well as tradeoffs inherent in seeking to optimize fairness along different metrics. As mentioned in the RMF, an organization must take a comprehensive approach to mitigating risk, while also optimizing operational effectiveness. Especially, the consideration of the risks, the appropriate way to define fairness, positive and negative impacts of *not* deploying AI tools, in comparison to deploying AI tools.

At Indeed, we believe in the fundamental benefit of appropriately developed and deployed AI tools to help connect job seekers and employers.

Functions, Categories, and Subcategories

We support the establishment of common definitions of key terms related to AI, and we support the use of plain language to enable the understanding and use of principles by a broad audience. Specifically, a stronger distinction between interpretability and explainability would benefit the RMF - including an example to illustrate the distinction between the two.

Alignment with Existing Risk Management Practices

The RMF places a large emphasis on up-front investment in risk identification (mapping) and documentation, which may not be readily adopted by development teams or smaller companies interested in innovation and advancement. A segmented framework, not only by application, but also by size of operation may benefit more stakeholders. As well, framework principles specific

to various applications may help in its use with data scientists and machine learning engineers who are used to exploratory work before committing to a modeling strategy.

Missing Information and Other Feedback

The RMF should include more actionable, implementable guidance. For example, what does scaling governance look like for segmented organizations? The variability in the tools, widespread applications, and constantly changing compliance requirements makes, in general, the usability of a broad framework more difficult to be applied and used.

As well, greater discussion and agreement should be made in the role of 'accountability.' In general, most stakeholders believe in appropriate accountability, Indeed included, however the RMF is unclear how accountability could be operationalized beyond regulation.

In addition, the relationship between 'fairness' and 'bias' requires further discussion and agreement. In its current form, the RMF uses a circular definition of 'fairness' and 'bias' by saying, 'fairness is the absence of bias.' Speaking specifically to the practical nature of mitigating bias, while seeking fairness, and retaining operational benefit to the users of a system needs further discussion. As noted above, the RMF could benefit from further engagement with the problem that different fairness metrics cannot be maximized simultaneously.

Conclusion

Thank you for the opportunity to provide feedback to NISTs AI Risk Framework: Initial Draft. We applaud NISTs efforts to create a framework which is both effective at mitigating risk and operational, which can be applied by a diverse set of stakeholders and users of AI tools and systems. Indeed stands ready to support NIST and offers to provide further background or perspectives as the framework is further developed. Please do not hesitate to reach out to Matthew Jensen at [...](#).