AI Risk Management Framework

National Institute of Standards and Technology

Feedback from Uptake on Initial Draft

INITIAL DRAFT PUBLISHED	Mar 17, 2022
FEEDBACK SUBMITTED	Apr 29, 2022

Contents

NIST UPTAKE

EXECUTIVE SUMMARY

Artificial intelligence (AI) has evolved from an experimental technology deployed in limited (and mostly government- or business-centric) use cases, to a pervasive component of the tools and technology used everyday by consumers and professionals. The evolution will undoubtedly continue at an ever-increasing pace.

As a leader in AI technologies for industrial and mobility use cases, Uptake commends the National Institute of Standards and Technology (NIST) for acknowledging the need for common language, standard frameworks, and shared best practices around the development, deployment, and management of AI systems – regardless of sector or scale. We believe NIST, as a non-regulatory organization with the broad directive of facilitating collaborative and effective innovation throughout the U.S. and our global community, is the appropriate steward for this effort.

The Initial Draft of NIST's AI Risk Management Framework (AI RMF) provides a sound foundation and roadmap for this effort. We strongly endorse NIST's clear effort to align the AI RMF with similar, existing frameworks; we appreciate NIST's emphasis on accessible, "plain English" terminology throughout the AI RMF; and we admire the innovative conceptual model provided by the AI RMF Core.

While we recognize that, at this early stage, the structure, language, and content of the AI RMF are in-progress, we identify the following high-impact areas for improvement.[1]

The AI RMF will benefit from a more cohesive flow between sections.

The utility of risk management frameworks, in general, is first-and-foremost born from the reader's ability to easily understand each concept and naturally draw extensions to their own organization or work products. We believe the AI RMF would be enhanced if it adopted a more user-friendly tone that's a bit less technical and formal.

Here are some examples of what we mean:

	Original Copy	**Modified Copy**
Page 14 Line 3-4	The AI RMF Core provides outcomes and actions that enable dialogue, understanding, and activities to manage AI risks.	The AI RMF Core is an organizing framework organizations may use to help stand up, evaluate, or enrich their AI risk management strategy.
Page 15 Line 13-15	Determination of whether AI use is appropriate or warranted can be considered in comparison to the status quo per a qualitative or more formal quantitative analysis of benefits, costs, and risks.	Qualitative and quantitative analyses of the benefits, costs, and risks of implementing an AI system can help determine whether the use case is trustworthy and ethical.

[1] For specific responses to the eight prompts NIST provided for this Initial Draft, please refer to page 4 of this submission.

We also believe the AI RMF would benefit from greater clarity when distinguisging between AI benefits and AI risks. In Section 5, the AI RMF defines a *trustworthy AI system* by providing the reader with knowledge of the principles of trustworthy AI and organizing frameworks to simplify the consumption and application of that knowledge. In parallel, the AI RMF seeks to also explain the relationship between *risky AI systems* and trustworthy AI systems. This may create confusion for non-technical or general audiences; separating the two may make for a simpler, more fluid alternative.

- We recommend sharpening the focus and language of Section 5 to focus more narrowly on defining the principles of trustworthy AI, and introducing a new Section (immediately following Section 5) focused on defining the characteristics of risky AI. Ideally, this new section will articulate the characteristics of risk through real-world, concrete examples, as well as relating back to the principles of trustworthy AI.

- Put simply, the AI RMF should define the ideal state (ie, trustworthy AI) and separately highlight scenarios or practices that diverge from it (ie, risky AI).

The AI RMF taxonomy requires more technical, AI-specific features.

The taxonomy must be reflective of both the principles of risk management and the nuances of AI systems. In its current state, the AI RMF taxonomy (Section 5) tilts heavily toward general risk management principles. While this is useful information, in order to better connect the RMF to the burgeoning AI world, we suggest aligning the risk management principles more closely and specifically to the world of AI.

- The enclosed graphic (AIRS AI Risk Categories – Appendix / page 6) demonstrates an alternative approach to classifying AI risk/trust. Notably, the AIRS approach employs language that aligns more with AI as a technical domain, less with risk management as a business practice. This would be a particularly useful approach for the recommended New Section on defining AI Risk (explained above).

- To be sure, NIST's focus on making the AI RMF a document that can be consumed by a diverse set of stakeholders – even those lacking technical expertise in AI – is critical to a successful outcome. We do not intend for this feedback to violate or diminish that.

* * *

Uptake is energized by the momentum NIST has ignited by not only committing to unravel the complexities of AI risk management, but engaging a broad spectrum of stakeholders throughout the process. The remainder of this submission includes the following:

§ 1. Direct responses to the eight areas of feedback specifically requested by NIST for the Initial Draft of the AI RMF. Beginning page 4.

§ 2. An appendix of supporting materials. Beginning page 6.

§ 3. Company overview of Uptake, provided as helpful context for Uptake's subject matter expertise and leadership in the Industrial AI space. Beginning page 7.

As future drafts of the AI RMF are developed, please know that Uptake will be on standby to collaborate further and more in-depth with NIST.

§ 1. DIRECT RESPONSES TO NIST'S REQUEST FOR FEEDBACK

1. **Whether the AI RMF appropriately covers and addresses AI risks, including with the right level of specificity for various use cases.**

   Discussed in Executive Summary.

2. **Whether the AI RMF is flexible enough to serve as a continuing resource considering evolving technology and standards landscape.**

   Yes, both the AI RMF Taxonomy and Core provide a degree of flexibility – primarily through the use of shared common language and an emphasis on concepts (theory) versus examples (utility) – that will allow the document to remain relevant without heavy modifications to content.

   As noted in the Executive Summary, however, we do believe future drafts of the AI RMF may benefit from incorporating more AI-specific tactical and technical detail – especially if done through language and terminology accessible to broad audiences.

3. **Whether the AI RMF enables decisions about how an organization can increase understanding of, communication about, and efforts to manage AI risks.**

   In its current state, we suspect the length and density of the AI RMF may deter general audiences (eg, non-technical, non-Risk Professionals). We discuss this point further in the Executive Summary.

   That aside, we do believe the Taxonomy and Core are incredibly valuable tools to help organizations understand the nuances of AI risk (via the Taxonomy, specifically) and evaluate the breadth and efficacy of its proactive AI risk management strategies (via the Core, specifically).

4. **Whether the functions, categories, and subcategories are complete, appropriate, and clearly stated.**

   In our view, the Core is the most compelling aspect of the AI RMF and provides an effective mechanism to help organizations scope and evaluate their planning and strategy for AI risk management.

   Where we find the Functions and Categories to cover an appropriate depth and breadth of detail, we see areas for improvement within the subcategories. Expanding the subcategories to provide more action-oriented guidance (or "links" out to recommended external resources) might make them not "analyze the risk" but "analyze the risk using X, Y, Z method") may help the subcategories be more complementary to their related Function / Category.

5. **Whether the AI RMF is in alignment with or leverages other frameworks and standards such as those developed or being developed by IEEE or ISO/IEC SC42.**

   Yes. The visual mapping of AI RMF Taxonomy to other existing frameworks (Fig 4, p9) is a very effective way of integrating "compatibility" within the document itself. To extend

this alignment even further, future drafts may consider incorporating commentary on the interplay between AI RMF and legacy frameworks like the Federal Reserve's guidance on financial model risk management (SR 11-7).

6. **Whether the AI RMF is in alignment with existing practices, and broader risk management practices.**

   Yes, the AI RMF provides an innovative and unique model, all the while remaining consistent with the basic principles and strategies of risk management. The AI RMF does a nice job of applying long-standing principles of risk management to the new world of AI.

7. **What might be missing from the AI RMF.**

   There are a few areas where the AI RMF may benefit from incorporating more language that aligns with AI as a technical domain. Discussed further in Executive Summary.

8. **Whether the soon to be published draft companion document citing AI risk management practices is useful as a complementary resource and what practices or standards should be added.**

   Yes, we believe the soon-to-be-published Practice Guide will be a valuable addition to the AI RMF. Expanding upon the feedback shared in #4, we see an opportunity for the practice guide (depending upon how it is developed) to fill the role currently played by the Core Subcategories.

§ 2. APPENDIX

The AIRS (AI Risk and Security) Working Group is a collective of 40+ practitioners and academics from varied backgrounds, including technology risk, information security, legal, privacy, architects, model risk management, and others, working for financial and technology organizations and academic institutions.

The figure at right shows an alternative approach to AI risk categorization and classification, developed by AIRS in their whitepaper on AI risk and governance.

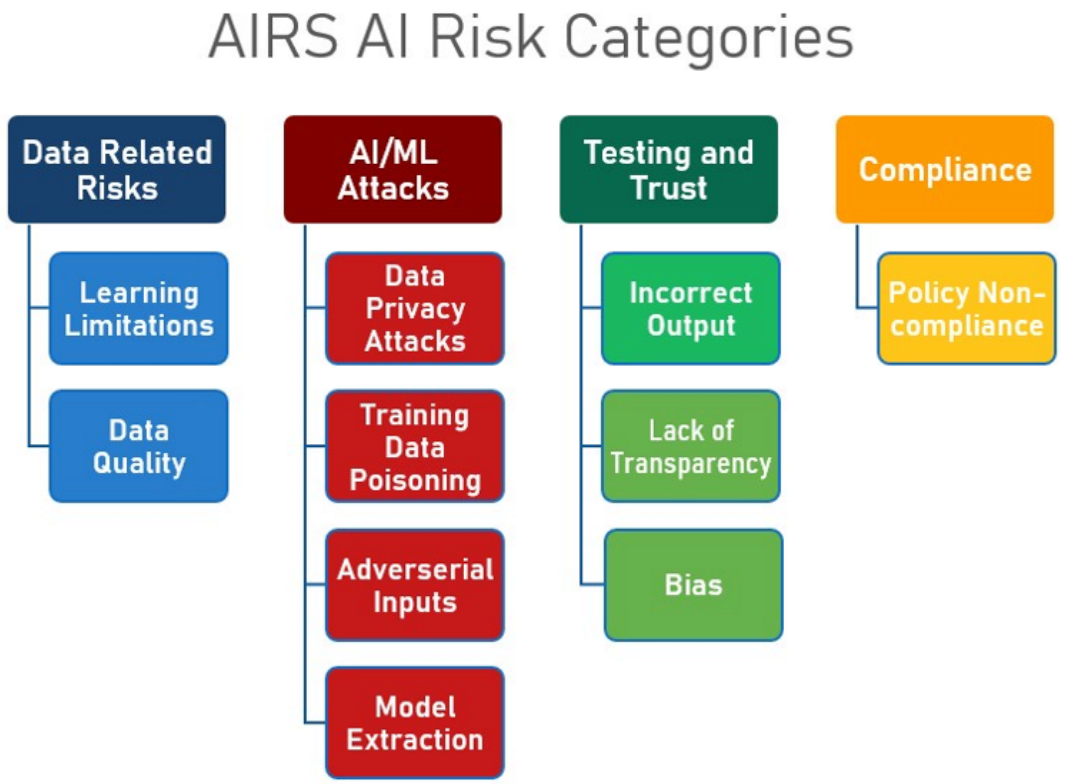

Image Source: ai.wharton.upenn.edu/wp-content/uploads/2020/12/AIRS-AI-Risk-Categories_Figure-1_Dec2020.png

AIRS Whitepaper: ai.wharton.upenn.edu/artificial-intelligence-risk-governance/

More info about AIRS: www.airsgroup.ai/

THE MOST TRUSTED SOURCE OF INDUSTRIAL INTELLIGENCE

READY TO UNLOCK YOUR DATA'S POTENTIAL?

info@uptake.com
uptake.com

For decades, industry leaders around the world saw big data as the destination; now they are realizing that it was just the beginning. Most industrial data was collected without a clear strategy or purpose, making the extraction of value a huge challenge. That's where Uptake comes in. Over the course of more than 100 customer engagements, Uptake has analyzed billions of hours of machine data from equipment across various industries. With Uptake, outcomes are achieved quickly, reliably, and at scale.

4M
predictions per hour

150B
data points processed

30+
patents issued

2,000+
analytics models

WHO IS UPTAKE?

Founded in 2014, Uptake is the intelligence system for industrial assets. Providing AI-driven asset performance management solutions, Uptake gives all departments — maintenance, reliability, and operations teams — a single, shared view of every piece of equipment in an operation. With the power of machine learning (ML) and artificial intelligence (AI), Uptake helps customers predict and prevent asset failures and unplanned downtime. Driven by powerful data science models, our software delivers insights that mitigate risk, optimize maintenance strategy and asset performance, reduce costs, and enhance safety. With 30+ patents issued, Uptake is recognized for leadership in the industrial internet of things (IIoT) by the World Economic Forum, CNBC, and Forbes.

THE UPTAKE DIFFERENCE

CONTENT

- Asset-specific advanced analytics bundles
- Asset Strategy Library® (ASL®)
- Value engineering

DATA SCIENCE

- World-class data science engines and data tools
- High-precision artificial intelligence and machine learning models
- Cross-industry network effects

EXPERTISE

- Digital domain proficiency across OT and IT
- Asset-intensive industries subject matter experts
- Library of high-value analytics use cases

SOFTWARE FOR HEAVY INDUSTRY

9%
availability improvement

15%
technician efficiency increase

20%
maintenance cost reduction

INTEGRATE AI ACROSS THE BUSINESS

Operationalizing AI cannot be achieved in isolation. Every insight produced by an analytics model is an opportunity to take action. The business process — fueled by the insight — is what delivers the business outcome. The Uptake Platform takes operationalizing and integrating AI just as seriously as the building and training of data science models, with powerful APIs and connectors to deliver insights to a customer's technology ecosystem. Rather than replacing existing applications and workflows, Uptake enhances current systems.

AUTOMATE DATA INTEGRITY

Data integrity is the most important ingredient in any successful AI initiative. It's far easier said than done. Data science teams typically spend up to 80% of their time gathering and wrangling data. Let Uptake do the heavy lifting. The Uptake Platform automates the complex task of ingesting and standardizing data to ensure data integrity. With Uptake, the data is always ready for advanced analytics.

GO FASTER WITH CURATED AND LEARNED CONTENT

Uptake has the world's largest library of industry knowledge with the Asset Strategy Library® (ASL®), over 2,000 pre-trained models for specific asset types, templatized asset data models, and curated content — all designed for the industrial sector. Created and fine-tuned based on data collected over billions of operating hours across more than a million assets, this library is the secret weapon to AI success.

AVOID TRIAL AND ERROR WITH PROVEN TECHNOLOGY AND DATA SCIENCE

As a fundamental component of the Uptake Platform, AI engines enable new models for common use cases to be configured in minutes. Models to predict machine failures, identify component survival probabilities, detect pattern anomalies, and prescribe recommended actions are just some of the analytics that can be put into production very quickly.

SENSOR DATA IS VALUABLE, BUT NOT REQUIRED

	ENGINEERING DATA	TRANSACTIONAL DATA	TIME SERIES DATA
Examples	• Design capabilities • Operational constraints	• Work orders	• Sensor readings • Fluid lab analyses • Derived calculations
Purpose	Provide context for analytics refinement	Maintenance program optimization (planning, financial & performance)	Real-time anomaly detection and failure prediction
Typical asset	All	All	Connected asset that is critical to operations
Industry focus	All	All	• Fleet • Wind energy • Mining • Manufacturing • Rail • Public Sector
Contribution to value *Based on potential value of a full rollout of Uptake's capabilities across a connected asset portfolio.*	0% Note: no incremental value, but required baseline for analytics	58%	42%

INDUSTRY ANALYST FEEDBACK

TRUSTED BY THE MOST REPUTABLE EXPERTS

ARC Advisory Group

"Uptake's overall strategy reflects the shift toward ready-to-deploy AI that leverages pre-trained industrial machine learning."

Gartner

"Uptake's comprehensive, cross-industry capabilities for Predictive Asset Management are competitive with other IIoT platforms."

[APM Market Guide 2020]

diginomica

"Uptake does more with IoT by 10 am most days than many ERP vendors get done in a year."

FORRESTER®

"The most striking thing about [Uptake's Platform] is the end-to-end thought-fullness of the solution."

constellation RESEARCH

"Uptake's approach is a ground-up rethinking of using [...] device data."

LNS research

"The culture and vision that has been developed [...] at Uptake may well be the most agile and adaptable of any we have seen."

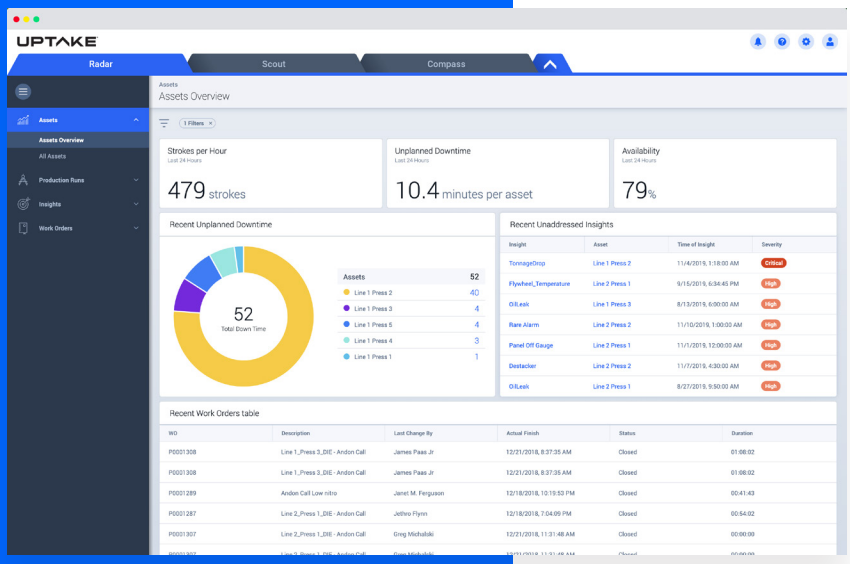

Avoid the trial and error. Uptake delivers immediate value from industrial analytics based on our work with industry leaders.

GLOBAL RECOGNITION

CRAIN'S CHICAGO BUSINESS MOST INNOVATIVE COMPANIES 2020

WORLD ECONOMIC FORUM Technology Pioneer

CNBC DISRUPTOR 50

Forbes AI 50

Gartner COOL VENDOR

CUSTOMER TESTIMONIALS

TANGIBLE VALUE ACROSS INDUSTRIES

"[What] Uptake does better than most is listen to product feedback and incorporate it into products."
— Fleet Analytics Manager

"[Uptake has] proven that AI and ML have moved beyond being buzzwords and can create tangible value."
— SVP of Equipment

"We have seen a 50% reduction in idle time since [we launched 6 months ago]."
— Senior Manager of Operations

"The application is simple to use, yet provides so much useful information."
— Sr. Manager Engineering & Sustainability

"This is one of the greatest things we've seen in a long time."
— Senior Fleet Manager

"Uptake's differentiation is the ability to create value almost immediately with [...] pre-trained models and industry expertise."
— Analytics Manager

"[Uptake's] constant development cycle and speed is much quicker than anyone else we work with."
— Fleet Analytics Manager

SELECT CUSTOMERS

FLEXCO

U.S.ARMY

Progress Rail
A Caterpillar Company

MIDAMERICAN ENERGY COMPANY

American Red Cross

BHE RENEWABLES

UnitedRoad

UNITED STATES MARINE CORPS

CATERPILLAR®