10. Effective ways to structure the Framework to achieve the desired goals, including, but not limited to, 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. Respondents are asked to identify any current models which would be effective. These could include – but are not limited to – the NIST Cybersecurity Framework or Privacy Framework, which focus on

| |
|---|
| outcomes, functions, categories and subcategories and also offer options for developing profiles reflecting current and desired approaches as well as tiers to describe degree of framework implementation; and |

These are personal views, and employers are not responsible for the views expressed here.

| **General RFI Topics** (Use as many lines as you like) | **Response #** | **Responding organization** | **Responder's name** |
|---|---|---|---|
| | | | |
| | | | |
| | | | |

| | | | |
|---|---|---|---|
| | | | |
| | | | |
| **Responses to Specific Request for information** (pages 11,12, 13 and 14 of the RFI) | NIST spells out its views about what features the AI RMF should have as: be consensus-driven, providing common definitions, using plain language, be adaptable, be risk-based, outcome-focused, voluntary, and non-prescriptive, be readily usable as part of any enterprise's broader risk management strategy, be a living document. While all of these are great suggestions, they would have been much more specific, had they been grounded in an AI Risk framework. Having a framework more specific to AI would allow for more structured collection of comments. Yet, that would be a iterative process. The current descriptions, though generic, are necessary scaffoldings for any good frameworks. It is good that NIST recognizes that AI Risk Framework is different to other generic risks. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |
| | | | |
| 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; | According to the OECD report (OECD, 2019) "AI has pervasive, far-reaching and global implications that are transforming societies, economic sectors and the world of work, and are likely to increasingly do so in the future."<br>Most practitioner's literature suggests that every risk management framework should consist of a minimum of risk identification, measurement, mitigation, reporting/ communication and monitoring, and governance. In addition, formulation of the risk problem and the context should also be included.<br>The challenge for AI based risk, as with any emerging technology, is that identification and assessment of risks are a moving target, and there is an element of projection required.<br>In addition, there are both ethical challenges such as biases in the AI, ethics of replacing human labour, removing large swaths of workforce from the market and AI making decisions inherently different from human decision making process, and so on. | Responder belongs to ARDC and UTS, but the views are that of the responder | |
| | | | |

| | | | |
|---|---|---|---|
| 2. How organizations currently define and manage characteristics of AI trustworthiness and whether there are important characteristics which should be considered in the Framework besides: accuracy, explainability and interpretability, reliability, privacy, robustness, safety, security (resilience), and mitigation of harmful bias, or harmful outcomes from misuse of the AI; | In the recent years, various governments (at least OECD) and jurisdictions have issued sets of guidelines for developing trustworthy AI systems. These are too numerous to name. Similarly, consulting organizations are claiming that they put trust at the centre of AI. The enterprises are just starting to realize ROI with AI applications, and the movement to make these systems ethical and responsible is still nascent.<br>That said, AI development has always included metrics such as accuracy and robustness. Now, the conversation has also shifted to explainability and interpretability as key dimensions needed to trust AI. Similarly, from a cybersecurity perspective security has become an important issue. Recent interest in ethical AI has resulted in a surge of interest in mitigation of harmful bias, misuse of AI as well as safety issues relating to AI.<br>One issue with trust along the ethical or bias dimension has been about ensuring whether AI systems include internal and external checks to help enable equitable application across all participants.<br><br>In this context, ARDC (and global counterpart Research Data Alliance) have a special role. The ARDC contributes to the digital global public good through its research infrastructure programs and its work on the FAIR Data principles with the Australian research community. These principles can be categorised as a global or open data standard. Though FAIR Data shares similarities with Open Data, the two concepts are not interchangeable. “Open Data” refers to “data that can be freely used, reused and redistributed to anyone - subject only, at most, to the requirement to attribute and share alike.” On the other hand, data is FAIR if it is “Findable,” “Accessible under well defined conditions,” “Interoperable” and “Reusable.” According to Mons et al, FAIR refers to a set of principles, focused on ensuring that research objects are reusable, and actually will be reused, and so become as valuable as possible … [These principles] describe characteristics and aspirations for systems and services to support the creation of valuable research outputs that could then be rigorously evaluated and extensively reused, with appropriate credit, to the benefit of both creator and user.<br><br>The FAIR Data principles were initially proposed in a 2016 paper authored by scholars working primarily in the life sciences. Though its origins reside within a field of scientific inquiry, advocates argue that the FAIR principles | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| | may be “equally applied to any data, or any service, in any discipline.” Since its introduction, the FAIR Data principles have been adopted by governments, universities, and funding organisations.<br><br>FAIR data is a cornerstone of trustable data, which is a prerequisite for trustable AI. Therefore, FAIR principles should be considered as an important metric for trustable AI. | | |
| | | | |

| | | | |
|---|---|---|---|
| 3. How organizations currently define and manage principles of AI trustworthiness and whether there are important principles which should be considered in the Framework besides: transparency, fairness, and accountability; | A state of the discussion on AI trustworthiness, comprising a wide array of attributes such as robustness, accuracy, fairness, explainability, and privacy, has been explored and presented by the National Academies of Sciences, Engineering, and Medicine (NASEM), where NIST is also co-sponsor. Therefore, it is not necessary to repeat the same, but for reference, the info is given here: https://vimeo.com/showcase/8327408<br>In addition to these, it is encouraged that AI development take into account two important principles, namely FAIR and CARE. The former is about making the data "Findable," "Accessible under well defined conditions," "Interoperable" and "Reusable", where focus on characteristics of data that will facilitate increased data sharing among entities while ignoring power differentials and historical contexts. However, an unintended consequence is that the emphasis on greater data sharing alone creates a tension for Indigenous Peoples. The CARE Principles for Indigenous Data Governance was designed to address this issue, and are "people and purpose-oriented, reflecting the crucial role of data in advancing Indigenous innovation and self-determination". CARE principles complement the existing FAIR principles. Further info can be found at: https://www.gida-global.org/care<br>ARDC as an organization has exceptional depth of skills and expertise in these areas and would be able to assist, if required. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |
| | | | |

| | | | |
|---|---|---|---|
| 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; | Why AI RMF is different:<br>Framework is different:<br>AI poses unfamiliar risks and creates new responsibilities<br>Over the past two years, AI has increasingly affected a wide range of risk types, including model, compliance, operational, legal, reputational, and regulatory risks. Many of these risks are new and unfamiliar in industries without a history of widespread analytics use and established model management. And even in industries that have a history of managing these risks, AI makes the risks manifest in new and challenging ways. For example, banks have long worried about bias among individual employees when providing consumer advice. But when employees are delivering advice based on AI recommendations, the risk is not that one piece of individual advice is biased but that, if the AI recommendations are biased, the institution is actually systematizing bias into the decision-making process. How the organization controls bias is very different in these two cases.<br>These additional risks also stand to tax risk-management teams that are already being stretched thin. For example, as companies grow more concerned about reputational risk, leaders are asking risk-management teams to govern a broader range of models and tools, supporting anything from marketing and internal business decisions to customer service. In industries with less defined risk governance, leaders will have to grapple with figuring out who should be responsible for identifying and managing AI risks.<br><br>AI is difficult to track across the enterprise<br>As AI has become more critical to driving performance and as user-friendly machine-learning software has become increasingly viable, AI use is becoming widespread and, in many institutions, decentralized across the enterprise, making it difficult for risk managers to track. Also, AI solutions are increasingly embedded in vendor-provided software, hardware, and software-enabled services deployed by individual business units, potentially introducing new, unchecked risks. A global product-sales organization, for example, might choose to take advantage of a new AI feature offered in a monthly update to their vendor-provided customer-relationship-management (CRM) package without realizing that it raises new and diverse data-privacy and compliance risks in several of their geographies. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| | Compounding the challenge is the fact that AI risks cut across traditional control areas—model, legal, data privacy, compliance, and reputational—that are often siloed and not well coordinated.<br>AI risk management involves many design choices for firms without an established risk-management function<br>Building capabilities in AI risk management from the ground up has its advantages but also poses challenges. Without a legacy structure to build upon, companies must make numerous design choices without a lot of internal expertise, while trying to build the capability rapidly. What level of MRM investment is appropriate, given the AI risk assessments across the portfolio of AI applications? Should reputational risk management for a a global organization be governed at headquarters or on a national basis? How should we combine AI risk management with the management of other risks, such as data privacy, cybersecurity, and data ethics? These are just a few of the many choices that organizations must make.<br>Models are different:<br>Model risk management (MRM) in regulated industries such as banking is currently performed by dedicated and independent teams reporting to the chief risk officer. While these firms have developed a robust MRM approach to improve the governance and control of their critical models determining capital requirements and lending decisions, this approach is usually not ideal for firms with different requirements or in less heavily regulated industries, for the following reasons:<br><br>MRM is typically based on a point-in-time model assessment (for example, once every one to five years), which assumes that the models are largely static between reviews. AI models learn from data, and their logic changes when they are retrained to learn from new data. For example, a fraud model is retrained weekly in order to adapt to new scams.<br>Traditional MRM workflows are often sequential and require six to 12 weeks of review time after the model development is complete, which delays deployment. These workflows are not easily adapted to the agile and iterative development cycles frequently used in AI model development.<br>MRM is often focused more on traditional risk types (primarily financial risks, such as capital adequacy and credit risk) and may not fully cover the | | |

| | | | |
|---|---|---|---|
| | new and more diverse risks arising from widespread use of AI such as reputational risk, consumer and conduct risk, and employee risk.<br>Some applications and use cases, such as chatbots, natural-language processing, and HR analytics, can qualify as “models” under regulatory definitions used in banking. But these applications are very different from the traditional model types (for example, capital models, stress-testing models, and credit-risk models), and traditional MRM approaches are not easily applied.<br>AI and machine-learning algorithms are often embedded in larger AI application systems, such as software-as-a-service (SaaS) offerings from vendors, in ways that are significantly more complex and more opaque than traditional models. This greatly complicates coordination between those who review the model and those who assess the application and platform (IT risk) or the vendor (third-party risk). | | |
| | | | |

| | | | |
|---|---|---|---|
| 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 above; | There are no frameworks for management of AI based risks. However, as we will explain later, there are frameworks in releated fields such as as well as cybersecurity, general risk etc that have some features to accommodate AI based risk.<br><br>We have compared the framework as shown in the later question. At a high (proverbial 50,000ft) level, all risk frameworks have comparable features by sharing similar cardinal activities of risk identification, measurement, mitigation, reporting/ communication and monitoring, and governance. A governance framework, while not being a risk management framework, also allow for similar activities.<br>Truth is that none of them are actually adequate to cover AI risks. Instead, organizations are managing risks through policies and rules by the seat of their pants. These are very reactive, responding to external legislations such as privacy and security. However, most frameworks allow for the above list of attributes to be consensus driven, adaptable, risk based, outcome focused, voluntary and so on.<br>NIST spells out its views about what features the AI RMF should have as: be concensus-driven, providing common definitions, using plain language, be adaptable, be risk-based, outcome-focused, voluntary, and non-prescriptive, be readily usable as part of any enterprise's broader risk management strategy, be a living document. While all of these are great suggestions, a specific conceptual model or reference framework would help ground the elicitation. Having a framework more specific to AI would allow for more structured collection of comments. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |
| | | | |

| | | | |
|---|---|---|---|
| 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; | There are multiple ways to govern potential risks and any risk management framework should be tailored to each contexts, including but not limited to the domain, the type of AI being developed, the data available, what or who is at stake and who is at risk and so on.<br><br>There are a number of risk management frameworks found in the literature, including those for financial risk, chemical process risk, nuclear risk, cyber security risk, health risk, construction or project risk, and so on. E.g. COSO ERM, ISO 31000 risk management guidelines, the AS/NZS 4360 standard and the Canadian Framework for the Management of Risk.<br><br>None of these are directly used for AI related risk management. For that matter, AI RMF is not yet an official area of execution. However, AI is managed through general data, IT, and cyber security governance frameworks. These may include definitions, inventory, policy/standards, and framework, including controls.<br><br>Obviously, NIST itself has a Cybersecurity Framework (Joint Task Force, 2018) that supports NIST's security and privacy risk management standards, guidelines, and practices. These examples include support for an Enterprise Risk Management (ERM) approach in alignment with OMB and FISMA requirements that agency heads "manage risk commensurate with the magnitude of harm that would result from unauthorized access, use, disclosure, disruption, modification, or destruction of a federal information system or federal information." Similarly, NIST also offers various guidelines, an example of which is a guide entitled "Managing the Security of Information Exchanges". This guide is about managing the information exchange lifecycle and process.<br><br>Joint Task Force (2018) Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-37, Rev. 2. https://doi.org/10.6028/NIST.SP.800-37r2 | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |
| | | | |

| | | | |
|---|---|---|---|
| | When looking for AI RMF solutions, existing regulatory and supervisory frameworks and standards are a good place start. According to National Science and Technology Council, AI-related risk "falls within the bounds of an existing regulatory regime, . . . the policy discussion should start by considering whether the existing regulations already adequately address the risk, or whether they need to be adapted to the addition of AI."12 A recent report by the U.S. Department of the Treasury reaches a similar conclusion with regard to financial services.<br><br>When designing, AI specific risk models, tools and instruments, it is important to consider all dimensions that AI touches an organization, namely: (i) adoption, (ii) application, (iii) business model (iv) workforce transformation, and (v) regulation.<br><br>As mentioned earlier, there are a number of risk management frameworks found in the literature, including those for financial risk, chemical process risk, nuclear risk, cyber security risk, health risk, construction or project risk, and so on. At a broader, organizational level, frameworks such as the COSO ERM, the ISO 31000 risk management guidelines, the AS/NZS 4360 standard and the framework for the Management of Risk – Canada exist (Raz & Hillson, 2005; Frigo & Anderson, 2014; Ahmad et al., 2014; Agarwal & Ansell, 2016) have been applied.<br>Each RMF/ standards is geared towards a different purpose, but each also provides a solution to a puzzle that AI Risk might encounter. The Committee of Sponsoring Organizations of the Treadway Commission (COSO) updated the ERM framework to promote better integration into business management. Owing to its business management (and less technical) role, COSO is abstract and conceptual. Some practioners claim that it lack in specificity sufficient to design implementation tools, audit tests and other instruments of actual relevance (Rubino & Vitolla, 2014). In general, they address phases of the risk process at an abstract level and leave out more specific aspects related to the management of processes and the definition of policies and procedures.<br>On the other hand, frameworks, standards and processes from technical disciplines tend to be more useful. For example, IT governance frameworks such as COBIT framework has control frameworks and expanding their reach (Ridley et al., 2004; Tuttle & Vandervelde, 2007; | | |
| 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; | | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| | Rubino & Vitolla, 2012b and 2014a).<br>Similarly, engineering disciplines such as aeronautical, military, chemical and nuclear hazards and risk assessment is a very mature field and has actual resources that can really be implemented, as opposed to just talked about. These disciplines also introduce the concept of systems safety, rather than addressing just at component or siloed levels.<br>However, there are some key limitations. The three traditional hazards in these industries, fire, explosion, and toxic release, attacks etc., have remained largely unchanged for many years and the assessment and management techniques reflect this stability. This is not a luxury that AI would offer.<br>Systems Engineering, and Systems Sciences disciplines bring a serious, integrated perspective to risk, that is hitherto ignored in financial and enterprise risks, and even under-represented in cybersecurity frameworks.<br>Most accident investigation and analysis rests on the use of event-chain models, i.e., the accident causation is described as a chain of failure events and human errors that led up to the actual loss event. Such models are limited in their ability to handle system accidents (arising from dysfunctional interactions among components and not just component failures), software-related accidents, complex human decision-making, and system adaptation or migration toward an accident over time.<br>In response to the limitation of event chain models, models based on systems theory have been proposed for use in accident analysis (see, for example Rasmussen (1997). STAMP (Systems-Theoretic Accident Modeling and Processes) is one such model that has been proposed (Levenson et.al., 2003).<br>In contrast to largely mono-causal linear models, system models treat accidents as being coproduced from the systems component interactions as well as deviant behaviour of the system. This approach is particularly useful for understanding as well as intervening in the system to reduce accident rates.<br>When comparing to mature industry risk models, one of the most important distinction for AI is that not all potential hazards and consequences are knowable now—the framework and tools should be continually vigilant for new issues in the rapidly evolving area of AI. Therefore, it is necessary to create a new tools and frameworks that would suit the AI case, but not from scratch. These existing tools can | | |

| | | | |
|---|---|---|---|
| | contribute to tools in the AI RMF domain.<br>From the practice side, most consulting organizations have discussed AI risk. Oftentimes, they propose a framework to address the risk. In most cases, this is simply a summary of cyber security risk. Nonetheless, McKinsey's AI risk management framework is well aligned to the AI/ ML life cycle, with stages such as (1) Ideation, (2) Data Acquisition, (3) Model Development and (4) Industrialization, Monitoring and Maintenance. There are cardinal questions in each stage.<br><br>For example, determination of bias, setting up a team to manage bias as well as regulatory guidance. It also assesses maturity at the start, at the ideation stage. Similarly data sourcing begs questions that helps define which data sets are "off-limits" (for example, because of personal-privacy considerations) and which bias tests are required. In the Model development, the focus is on the selection of the appropriate method for transparency and interpretability. In order to manage risks in productionization, the framework recommend to define the monitoring requirements. While these are systematic, it essentially consists of two risks talked about in the AI field, namely biases and productionization. In that sense, the framework is not novel, but is a good starting point.<br><br>https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/derisking-ai-by-design-how-to-build-risk-management-into-ai-development | | |
| | | | |
| | | | |

| | | | |
|---|---|---|---|
| 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. | Measuring costs and benefits, and assessing the issues relating to inclusiveness is just one of the many issues plaguing AI risk. There is belief that cognitive diversity in AI teams would bring considerable level of controls or self-governance. It is also that inclusiveness, apart from its moral stance, could help rectify concerns pertaining to included stakeholders.<br>Some start-ups are already developing tools for managing AI Risk, but early versions are likely to be a version of compliance tools or checklists. For example, Fairly promotes itself as AI Governance, Risk, and Control (GRC) solution.<br>The WHO has recently published a guidance on Ethics & Governance of Artificial Intelligence for Health. It states that AI already helps to quickly diagnose disease, assists with clinical care, strengthens research and drug development, and supports public health interventions such as outbreak responses. But it also cautions against overestimating AI's benefits in health and the unethical gathering and use of health data, biases in algorithms and risks to patient safety, cybersecurity and the environment. The WHO report provides six principles for use in regulation and governance of AI in health:<br><br>Protecting human autonomy, Promoting human well-being and safety and the public interest, Ensuring transparency, explainability and intelligibility, Fostering responsibility and accountability, Ensuring inclusiveness and equity, and Promoting AI that is responsive and sustainable.<br><br>That said, those who participate in the AI workforce are increasingly getting skewed in distribution. Some academics believe that; 'technological progress is going to leave behind some people, perhaps even a lot of people, as it races ahead. There has never been a worse time to be a worker with only 'ordinary' skills and abilities to offer (Brynjolfsson, Ford and McAfee. 2012).'<br>At present, the rules that preside over AI technology revolve around the idea that whoever developed the technology, has complete autonomy and control over how it can be used and where it can be implemented. If we wish to ensure that AI technology is distributed fairly amongst societies, we must begin seeing AI systems not as private commodities, but instead as a potential benefit that can improve the quality of life for | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| | all of humanity. For such perspectives to prevail, it is fundamental that government agencies around the world begin developing and implementing national public policies that support such initiatives. By implementing universal measures as such, the international arena is unifying as one when it declares AI technology as a humanistic right, regardless of one's location, race, culture, nationality, colour or social class.<br>Recently, given the court ruling that AIs cannot patent their inventions, so whatever an AI output is, it's not going to be easy to legally control. | | |
| | | | |
| 9. The appropriateness of the attributes NIST has developed for the AI Risk Management | NIST spells out its views about what features the AI RMF should have as: be concensus-driven, providing common definitjions, using plain language, be adaptable, be risk-based, outcome-focused, voluntary, and non-prescriptive, be readily usable as part of any enterprise's broader risk management strategy, be a living document. While all of these are great suggestions, being grounded in a specific framework would be better for ideass elicitations. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| Framework. (See above, “AI RMF Development and Attributes”); | | | |
| | | | |
| 10. Effective ways to structure the Framework to achieve the desired goals, including, but not limited to, 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. Respondents are asked to identify any current models which would be effective. These could include – | We need to establish a baseline (table stakes) for AI risk management. This would include any technology risk management with associated frameworks. The existing approach to risk has been well discussed in Aven (2006). There is no need to reinvent the wheel there. As we know, the risk should also be integrated with organizational process. Systemic risk models which examined the idea that systems failures, rather than just human failure, were a major contributor to accidents (Hollnagel, 2004) began to address some of these issues (but not non-linear concepts) and recognised that events do not happen in isolation of the systemic environment in which they occur.<br>There are two key factors that often determine risk: technology factors with inherent nature of risk, and the human and organizational factors that often makes it worse. What Leveson realized is that when complexity is high within a system. This is because will be in social or socio-technical systems such as AI, the traditional approach loses its effectiveness. Even when components are seemingly risk free, the interactions could create new risk, and increase the risk levels significantly. Typically, in risk, machines are more reliable and are protected from human operators. With a seemingly purposeful entity such as AI, this may not be the case. There are always multiple goals and constraints for any organizational system — the challenge in design and risk management of AI is to identify and analyze the conflicts, to make appropriate tradeoffs among the conflicting requirements and constraints, and to find ways to increase system safety without decreasing system reliability.<br>Siloe in the organizations may not just lead to sub-optimal performance, but also increase the risk in complex sociotechnical systems where AIs are built and embedded. Each local siloes may act optimally in their limited context but together they could perpetuate larger risks in dysfunctional ways.<br>Therefore, the current frameworks, systems and tools would be far from | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| but are not limited to – the NIST Cybersecurity Framework or Privacy Framework, which focus on outcomes, functions, categories and subcategories and also offer options for developing profiles reflecting current and desired approaches as well as tiers to describe degree of framework implementation; and | sufficient to keep up with the amount of complexity we want to build into AI.<br>With respect to risk, we often talk about black swans to describe events that could not be predicted or fathomed. In reality, three types of swans need to be identified – white, grey and black swans – and handled differently, even though they fall on the continuum of uncertainty. They cannot be mixed, as they have completely different mathematical and statistical characteristics (Barta & Görcsi, 2021). Most risk management processes coming from accounting, finance or legal are driven towards compliance, trying to tackle the white swans mostly, and perhaps some well-known grey swans.<br>AI will spawn all three categories. Provisions must be made to identify, assess, and manage these in an integrated manner.<br>Aven T (2006) Risk assessment and risk management: review of recent advances on their foundation. Eur. J. Oper. Res., 253 (2016), pp. 1-13, 10.1177/1748006X17699145<br>Barta, G., & Görcsi, G. (2021). Risk management considerations for artificial intelligence business applications. International Journal of Economics and Business Research, 21(1), 87–106. https://doi.org/10.1504/IJEBR.2021.112012 | | |
| | | | |

| | | | |
|---|---|---|---|
| 11. How the Framework could be developed to advance the recruitment, hiring, development, and retention of a knowledgeable and skilled workforce necessary to perform AI-related functions within organizations. | Firstly recruitment, hiring, development, and retention of AI related workforce is occurs through a 1-dimensional approach at present, and the trend is only getting acute and worse.<br>The industry appears to recruit two types of talents:<br>Work-horses/ Doers: Those who can actually develop and deliver AI are recruited almost entirely from talented pool of coders (programmers) and the AI problem is pushed to become an engineering problem. Here, the effort is being made to avoid recruiting those with broader skillsets.<br>Leaders: There have been some attempts to move leaders laterally into AI domain. Here, effort is being made to avoid those with technical skills, but those who have the necessary leadership skills.<br><br>This has created a chasm between leaders and do-ers, and this has long standing implications.<br>Firstly, there is a need for mentorship roles in AI development. A mentor is someone who can direct the technical team, but also could speak the language of the leadership, and translate the requirements or vision.<br>Secondly, the team of doers need to be a broader set. At the very least, one could separate the skills based on the stage of the AI project lifecycle: Ideators, data wranglers, model developers, deployment specialists.<br>Many of the core technical skills required for AI development are familiar to practitioners, researchers and personnel working with AI. Some elements that are missing: Of particular importance is recognising the role of findable, accessible, interoperable and reusable (FAIR) data and software skills to address the disparate nature of data, especially interoperability of data across disciplines. The key to sustainability is recognising the proper investment in communities and AI excellence. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |
| | | | |
| 12. The extent to which the Framework should include governance issues, including but not limited to make up of design and | The governance issues are inextricably linked to risk. The frameworks also are related. A good example is the AI Governance Framework developed by Singapore. In January 2019, Singapore's Model AI Governance Framework has addressed ethical issues in a practical manner, and could be employed as part of AI Risk Framework. The framework considers internal governance structures and measures, determining the level of human involvement in AI-Augmented Decision-Making, Operations Management and Stakeholder Interactions and Communications. | Responder belongs to ARDC and UTS, but the views are that of the responder | Gnana K Bharathy |

| | | | |
|---|---|---|---|
| development teams, monitoring and evaluation, and grievance and redress. | | | |