

RFI Response

September 15, 2021

Virginia Tech Applied Research Corporation (VT-ARC)

In Partnership with

**Virginia Polytechnic Institute and University (VT) &
Commonwealth Cyber Initiative (CCI)**

(Comments included from R. Kuhn, R. Kacker, M.S. Raunak, NIST)

Contract Name: Artificial Intelligence Risk Management Framework

Notice ID: NIST-2021-0004; Docket Number: 210726-0151

Technical points of contact

Name – Laura Freeman; Feras A. Batarseh

Address - 900 N Glebe Rd, Arlington, VA 22203

Telephone Number - (571) 384-3833

Facsimile Number - (571) 449-3713

E-mail Address – laura.freeman@vt.edu; batarseh@vt.edu

VT-ARC Organizational DUNS: 078432677

VT Organizational DUNS: 00-313-7015

Disclosure of Information Statement:

This white paper includes data that shall not be disclosed outside the Government, except to non-Government personnel for evaluation purposes, and shall not be duplicated, used, or disclosed -- in whole or in part -- for any purpose other than to evaluate this submission. If, however, an agreement is issued to this Company as a result of -- or in connection with -- the submission of this data, the Government shall have the right to duplicate, use, or disclose the data to the extent agreed upon by both parties in the resulting agreement. This restriction does not limit the Government's right to use information contained in this data if it is obtained from another source without restriction.

The data subject to this restriction are contained in **ALL** sheets contained within.

A. Statement of Interest

Both VT-ARC and VT are very interested in supporting the National Institute of Standards and Technology (NIST) with development of a framework that can be used to improve the management of risks to individuals, organizations, and society associated with artificial intelligence (AI). VT-ARC and VT are integrated in their efforts supporting many different Artificial Intelligence (AI) and Machine Learning (ML) programs and provide the most cutting-edge technical advances in the AI/ML area.

Virginia Tech Applied Research Corporation (VT-ARC). VT-ARC is a private non-profit 501(c)(3), applied research corporation affiliated with Virginia Polytechnic Institute and State University (Virginia Tech or VT). VT-ARC accelerates solutions to complex issues of national importance through the identification, development and application of innovative analytics and advanced technologies. We leverage the rich, multidisciplinary research and innovation ecosystem of Virginia Tech, combining strategy, policy, technology, and operational considerations across multiple domains. VT-ARC has past performance over the past 10 years with multiple federal government organizations, industry partners, DOD innovation organizations, and other higher education and research institutions. VT-ARC has performed on two separate PIAs, one with ARL for nearly five years, and the other with AFRL/AFOSR for seven years. Among our partners are the ARL, the ARO, AFOSR, OUSDR&E, and DTRA. We currently support 16 research, development, planning and testing programs.

VT-ARC's headquarters, located in Arlington, VA, has over 16,000 total square feet of space, including 13,000 square feet of unclassified office space, additional TS secure classified spaces, and building access to venues that can support events up to 200 people. Our second office is located in Blacksburg, VA. While VT-ARC doesn't possess designated laboratory space, we can access university laboratory space through Virginia Tech.

VT Hume Center Intelligent Systems Lab (ISL). The Virginia Tech Hume Center believes providing research opportunities to students is key for developing the next generation of national security and technology leaders. Our research opportunities allow students to gain hands-on experience in their preferred discipline, network and build connections with award-winning and expert-level researchers in the field and explore an academic discipline more fully to develop interests into passions and careers.

The lines between data science, machine learning, and cybersecurity are blurring as algorithms that include machine learning are integrated into production systems. The validation of complex systems through rigorous test and evaluation processes is also needed to ensure the efficacy and safety of algorithms embedded in systems as they accomplish tasks with greater autonomy and operational impact. The design and development of this systems need to reflect their intended operating environment, representative human users, and the operational mission/tasks. The Hume Center's ISL conducts research to address critical areas of national security in three technological thrusts: 1) data science, machine learning, artificial intelligence, 2) cybersecurity and complex systems engineering, and 3) complex system design, validation, and test and evaluation (T&E).

Commonwealth Cyber Initiative (CCI). VT has a leadership role in developing AI assurance methods, and in managing the CCI AI Testbed. CCI supports over 70 participants from over 30 different Virginia universities and colleges. The AI assurance team at CCI develops model

agnostic and model specific assurance pipelines for a variety of applications including for public policy, government systems, American critical infrastructure (such as smart grids and smart farms), and other domains. The infrastructure within the AI testbed enables the widescale testing of such deployments, including clusters designed to allocate logically separated work environments. Containers that enable users to run their data science, machine learning (ML), and/or AI workloads. This construct enables different modes of operation and the ingest of relevant data to support threat detection, identification, framework alignment, correlation, and analytics workflows. These workflows can be staged on the testbed where various types of data science techniques can be applied. The research and development main focus of the AI assurance team at CCI is the assurance and risk mitigation of AI systems within different domains and sectors of society. The team is currently spearheading a publication of a book (to be released on March, 2022) on the assurance of AI systems and methods for AI risk management. This team is very excited to support this request from the NIST.

NIST REQUESTED INFORMATION

VT and VT-ARC have reviewed the requested NIST RFI and AI Framework. Two of our PhDs in the VT Hume Intelligence Systems Lab, Dr. Laura Freeman and Dr. Feras A. Batarseh are well published in areas paralleling your AI Framework effort. Many of the NIST AI Framework characteristics are specifically called out and explained in select recent manuscripts. The first paper: a survey on AI assurance [1] provides a roadmap to challenges that need to be resolved to create effective risk management frameworks. The second paper: enabling AI adoption through assurance [2], proposes a generic framework that addresses AI risk challenges. The third paper provides metrics for testing machine learning algorithms [3]. Test and Evaluation for AI [4] argues that a test and evaluation process where metrics as discussed in [3] are captured is key to a risk-based framework for supporting the decision to deploy AI into defense systems. The final reference provides a first pass at critical constructs to be developed in a framework for assuring autonomous intelligent agents. Incorporating some of the terminology and framework characteristics should be beneficial to your final proposal. Our team has taken excerpts from these documents and commented in the appropriate sections of the RFI. The five documents are referenced up front and attached in this .pdf response.

We look forward to supporting you on this effort!

1. Batarseh, F., Freeman, L. and Huang, C., "A Survey on Artificial Intelligence Assurance", *Journal of Big Data* 8, 60, Apr 2021.
2. Freeman, L., Rahman, A., Batarseh, F., "Enabling Artificial Intelligence Adoption Through Assurance", accepted at MDPI's *Journal of Social Sciences, Artificial Intelligence for Policy Analysis, Governance and Society (AI-PAGES)*, Sep 2021.
3. Lanus, E., Freeman, L., Kuhn, R. and Kacker, R. "Combinatorial Testing Metrics for Machine Learning," 2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW), 2021, pp. 81-84, doi: 10.1109/ICSTW52544.2021.00025.
4. Freeman, L. (2020). Test and Evaluation for Artificial Intelligence. *INSIGHT*, 23(1), 27-30.
5. Lanus, Erin, et al. "Test and Evaluation Framework for Multi-Agent Systems of Autonomous Intelligent Agents." *arXiv preprint arXiv:2101.10430* (2021).

NIST is requesting information related to the following topics:

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

VT/VT-ARC Comments: Another way of stating how an agency can manage AI risks is through evaluation of the AI system stated as assurance, validation, verification. [1]. The areas of focus for management and organization of AI are domain, AI subarea and the AI goal. [1]

An overall definition of AI assurance would add significantly to the user’s focus when analyzing what NIST is trying to accomplish and how NIST is going to manage AI risk. Based on prior definitions and recent AI challenges, we propose the following definition for AI assurance:

A process that is applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy, and explainable to a layman, ethical in the context of its deployment, unbiased in its learning, and fair to its users. [1]

Areas to assist with management of AI and essential to all forms of AI Assurance

- (1) Context: refers to the scope of the system, which could be associated with a timeframe, a geographical area, specific set of users, and any other system environmental specifications
 - (2) Correlation: the amount of relevance between the variables, this is usually part of exploratory analysis, however, it is key to understand which dependent variables are correlated and which ones are not
 - (3) Causation: the study of cause and effect, i.e., which variables directly cause the outcome to change (increase or decrease) in any fashion
 - (4) Distribution: whether a normal distribution is assumed or not. Data distribution of the inputted dependent variables can dictate which models are best suited for the problem at hand
 - (5) Attribution: aims at allocating the variables in the dataset that have the strongest influence on the outcomes of the AI algorithm. [1]
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.

VT/VT-ARC Comments. An alternate AI Assurance definition is the probability that a system leveraging AI algorithms functions only as intended and is free of vulnerabilities throughout the life cycle of the system. High levels of assurance stem from all the planned and systematic activities applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy, and explainable to a layman, ethical in the context of its deployment, unbiased in its learning, and fair to its

stakeholders. The probability that a system functions as intended includes assessments across constructs to include trustworthiness, explainability, fairness, assessment of unintended biases, and the assessment of ethical implications. [2]. **Our recommended verbiage for this section is inclusion of the terms fairness, and ethical implications.**

The areas of focus for management and organization of AI are domain, AI subarea and the AI goal. [1]

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.

VT/VT-ARC Comments. In order to add considerations to this section, AI trustworthiness should first be defined. Trustworthiness includes concepts such as integrity, human-centered development and use, respect the law, convey that humans must be able to trust the AI algorithms for wide-scale (agency level) adoption. [2]

Additional principles contributing to AI trustworthiness include reliability, robustness, security, and ethical implications.

Each of these terms is defined below for consideration.

Reliability, Robustness - highlight the need to understand the consistency of algorithm prediction across varying domains and understand when they may fail

Secure - Algorithms and data should be protected and controlled at the appropriate level, algorithms need to be robust to adversarial action

Ethical - as AI scales solutions, it also scales mistakes, discrimination, and potential non-ethical outcomes. Algorithms ought to be assured to ensure ethical standards (within the context it is applied) are met [2]

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.

VT/VT-ARC Comments. Our team recommends adding the language - *Algorithms and data should be protected and controlled at the appropriate level and algorithms need to be robust to adversarial action.* [2]

In order to apply appropriate security protocols, our team recommends the following approach as stated in the RMF framework: Categorize the Information System, Select Security Controls, Implement Security Controls, Assess Security Controls, Authorize Information System, Monitor Security Controls

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

VT/VT-ARC Comments. AI measurement and risk assessment should stem from the problem definition. Some overall measurement areas are algorithm performance, bias/fairness, security, safety, trustworthiness, explainability, ethicality. Additionally, traditional measurements of systems should be considered to include system performance, utility, efficacy, usability, resilience, reliability, and maintainability. [2]

It is essential to evaluate the AI system's ability to respond correctly to inputs. One aspect of this is ensuring that the input model used for training and testing AI systems accurately reflects the range of inputs that will be encountered in use. We have developed methods for measuring input space coverage that can be applied to the problem of quantification in measurement areas of security, safety, and explainability [3].

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.

Today, the strongest regulatory requirements for software and systems safety are those used in commercial aviation, based on recommendations from the Radio Technical Commission for Aeronautics (RTCA), primarily RTCA DO-178C. The key method for testing life-critical software in aviation, as required in DO-178C, is the *modified condition/decision coverage* (MCDC) criterion, which requires that every decision within code takes every possible outcome, each condition within each decision takes every possible outcome, and every condition in every decision has been shown to independently affect the decision outcome.

The MCDC criterion is extraordinarily time-consuming and expensive to achieve, typically consuming 85% to 90% or more of the software development budget (NASA study). For autonomous systems, there is an even more significant assurance problem: large parts of these systems use neural network or other black-box software that is difficult if not impossible to verify and test according to the MCDC criterion, or other structural coverage measures such as branch or statement coverage. The reason is that the behavior of neural networks depends on connections formed based on large volumes of input, and not on hard-coded logic and decision predicates as in conventional software. Testing and assurance of autonomous systems including neural nets is an unsolved problem and the subject of much current research. Methods of addressing assurance in this space have focused on measuring neuron coverage, but it is far from clear if neuron coverage has a significant relationship with correctness and safety in autonomous systems. It may be that neuron coverage is too simple of a metric that is not sufficiently indicative of system behavior, just as statement coverage in conventional software is a poor measure of test thoroughness or correctness. Most efforts at assuring correct operation of autonomous systems have been based on testing for extended periods, with inadequate measures of how thorough such testing has been other than time, number of miles driven, or other basic quantities. Such brute force testing may be insufficient for preventing failures under extremely rare circumstances. A well-known example is the fatal crash of a car in autonomous mode that resulted from a very rare four-factor combination of a white truck against a brightly lit sky, along with truck height and angle versus the car [1,2]. Measurements of input space model coverage are one approach to addressing the problem of

ensuring that inputs to the AI system are processed correctly, and these measures could be considered for future standards.

The IEEE Reliability Society is currently in the process of defining standards for autonomous systems. Notable with respect to trust and assurance in particular are IEEE 7001 and 7009, which address AI explainability and safety. These developing standards address a number of areas that are included in the proposed risk management framework.

- IEEE 7001 - Transparency of Autonomous Systems
 - <http://standards.ieee.org/develop/project/7001.html>
- IEEE 7009 - Fail-Safe Design - Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems
 - <https://standards.ieee.org/project/7009.html>

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

VT/VT-ARC Comments. The language ‘ensure that the AI RMF aligns with and supports other efforts’ is key to AI risk management. In order to adequately align with other efforts, the ‘Define’ phase of the process ensures all stakeholders are involved in the process. Stakeholders may include the mission/task champion (leadership), program management, system engineer, AI developer, requirements representative, test and evaluation personnel, end users, etc. depending on the application. The three most important elements in the ‘Define’ phase are the AI subarea, AI Domain, and the AI goals. Critical elements of this phase are:

AI Domain - what is the broad organizational mission that is acquiring the AI enabled system (e.g., government, energy sector, defense, healthcare, etc.)? What context does that bring to the challenge of assuring AI?

Mission context - what is the problem that the organization is trying to solve with an AI system(s)? How does the AI support the organizational mission?

AI Subarea - what is the type of AI needed (e.g., deep learning, reinforcement learning, etc.)? How does the AI contribute in ways that previous methods failed?

Scientific/Engineering Alignment - What are the scientific and/or engineering needs that the AI can solve? How does it integrate with known constraints? How does it incorporate prior knowledge (solely through data or other mechanisms)?

AI Goal -What are the primary directives for the AI, how does this stem from the AI domain? Must it be ethical, explainable, fair, etc.? [2]

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.

VT/VT-ARC Comments. Defining testing and assurance are critical to this section. Defining both of those here would be very helpful to the overall framework. We provide definitions applicable to all AI domains and subareas.

Testing: according to the American Software testing Qualification Board, testing is “the process consisting of all lifecycle activities, both static and dynamic, concerned with planning, preparation and evaluation of software products and related work products to determine that they satisfy specified requirements, to demonstrate that they are fit for purpose and to detect defects”. Based on that (and other reviewed definitions), testing includes both verification and validation. [2]

Verification is the “process of determining that a model implementation accurately represents the developer’s conceptual descriptions and specifications” [2]

Validation is the process of “determining the degree to which a model is an accurate representation.” [2]

Assurance: this term has been rarely applied to conventional software engineering; rather, it is used in the context of AI and learning algorithms. Our definition of AI assurance is included in our comments to section 1 of this document. Based on prior definitions and recent AI challenges, we propose the following definition for AI assurance:

A process that is applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy and explainable to a layman, ethical in the context of its deployment, unbiased in its learning, and fair to its stakeholders. [1]

9. The appropriateness of the attributes NIST has developed for the AI Risk Management Framework. (See above, “AI RMF Development and Attributes”)
 - Be consensus-driven and developed and regularly updated through an open, transparent process. All stakeholders should have the opportunity to contribute to the Framework’s development. NIST has a long track record of successfully and collaboratively working with a range of stakeholders to develop standards and guidelines. NIST will model its approach on the open, transparent, and collaborative approaches used to develop the Framework for Improving Critical Infrastructure Cybersecurity (“Cybersecurity Framework”) 4 as well as the Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management (“Privacy Framework”).5
 - Provide common definitions. The Framework should provide definitions and characterizations for aspects of AI risk and trustworthiness that are common and relevant across all sectors. The Framework should establish common AI risk taxonomy, terminology, and agreed-upon definitions, including that of trust and trustworthiness.
 - Use plain language that is understandable by a broad audience, including senior executives and those who are not AI professionals, while still of sufficient technical depth to be useful to practitioners across many domains.
 - Be adaptable to many different organizations, AI technologies, lifecycle phases, sectors, and uses. The Framework should be scalable to organizations of all sizes, public or private, in any sector, and operating within or across domestic borders. It

should be platform- and technology agnostic and customizable. It should meet the needs of AI designers, developers, users, and evaluators alike.

- Be risk-based, outcome-focused, voluntary, and non-prescriptive. The Framework should focus on the value of trustworthiness and related needs, capabilities, and outcomes. It should provide a catalog of outcomes and approaches to be used voluntarily, rather than a set of one-size-fits-all requirements, in order to: Foster innovation in design, development, use and evaluation of trustworthy and responsible AI systems; inform education and workforce development; and promote research on and adoption of effective solutions. The Framework should assist those designing, developing, using, and evaluating AI to better manage AI risks for their intended use cases or scenarios.
- Be readily usable as part of any enterprise's broader risk management strategy and processes.
- Be consistent, to the extent possible, with other approaches to managing AI risk. The Framework should, when possible, take advantage of and provide greater awareness of existing standards, guidelines, best practices, methodologies, and tools for managing AI risks whether presented as frameworks or in other formats. It should be law- and regulation-agnostic to support organizations' ability to operate under applicable domestic and international legal or regulatory regimes.
- Be a living document. The Framework should be capable of being readily updated as technology, understanding, and approaches to AI trustworthiness and uses of AI change and as stakeholders learn from implementing AI risk management.

VT/VT-ARC Comments. Our team concurs with the RMF Framework as presented and offers language for consideration to add additional context for all of the sections in this framework. This RMF framework meets our characteristics for AI assurance as reliable, dependable, explainable, and fair. AI assurance provides the necessary tools to enable AI adoption into applications, software, hardware, and complex systems. AI assurance involves quantifying capabilities and associating risks across deployments including: data quality to include inherent biases, algorithm performance, statistical errors, and algorithm trustworthiness and security. Our recommendation acknowledges that data, algorithmic, and context/domain-specific factors may change over time and impact the ability of AI systems in delivering accurate outcomes. [1]

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

Recommend adding this process from the International Statistical Engineering Association (Figure 1)

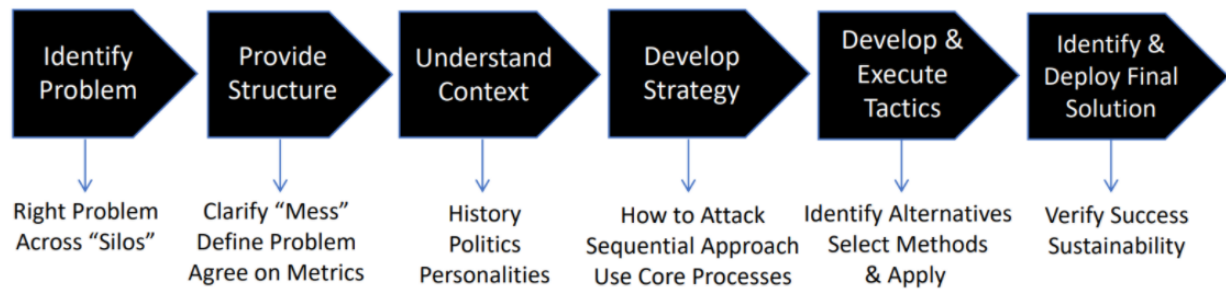


Figure 1. Typical Statistical Engineering Process, adapted from Hoerl and Snee (2017)[2]

Our overall process for AI assurance is six steps--Define, Measure, Characterize & Add Context, Plan Strategy, Execute & Analyze, and Monitor & Improve. [2]

Define. The define stage sets up the foundation for the AI assurance process. This phase should focus on ensuring that the AI assurance program is adequate. A key aspect of the define phases is ensuring that all stakeholders are involved in the defining process. Stakeholders may include the mission/task champion (leadership), program management, system engineer, AI developer, requirements representative, test and evaluation personnel, end users, etc. depending on the application. Based on an extensive literature review Batarseh, Freeman, and Huang [10] identify three elements of AI assurance that are important in the "Define" phase: AI Subarea, AI Domain, and AI Goals. [2]

Measure. There are numerous dimensions of AI measurement and not all intelligent systems will require all dimensions. The measures should stem from the problem definition. Dimensions of measurements that should be considered include: Algorithm Performance, Bias/Fairness, Security, Safety, Trustworthiness, Explainability, Ethicality. Although those measures are highly subjective, methods to quantify them are needed to further advancing the field. The loading of ethical standards or trustworthiness for instance into an AI system is -in most cases- domain specific; for instance, an ethical metric in a system built for a healthcare application will be -most likely- different than one in a warfare application, therefore, the process of "values loading" into the AI system is needed to capture those measures and create a benchmark to compare against for assurance quantification. [2]

Characterize & Add Context. The measurement of intelligent systems leveraging AI is dependent on the multi-dimensional space that includes the operational environment, system, hardware infrastructure, and other factors effecting the AI. The concept of an operating envelope applies at each view level. For example at the operational environment view one might be concerned about factors that include environmental conditions, operational users, mission tasks, and other factors external to the system impact performance outcomes. At the lower level the operating envelope many consist of data views like data quality, information content, domain, and range. Each of these views needs to be considered for their impact on achieving the goals of the intelligent system and appropriately planned for in the execution of the assurance program. [2]

Plan Strategy. We emphasize the need for leveraging test and evaluation across the system development, deployment, and operations life-cycle as the basis for a strategy for assuring AI. Our research highlights the need for an iterative processes and several quality improvement tasks that include multiple types of testing, but also debugging, manual inspection, and static analysis. While formal verification strategies can be a component of an assurance strategy, the complexity of statistical learning algorithms - especially when considered as employed as part of a system that interacts with the operational environment - demands a data driven approach to assurance. The planning strategy should consider the use of formal verification approaches as well as test venues that enable the exploration of AI model outputs directly, the incorporation of digital simulations, software-in-the-loop simulations, hardware-in-the loops simulations, full system testing in controlled environments, and full system testing in operational environments with operational users. The planning strategy should describe the iterative, sequential process for gathering knowledge across the key measurements for AI assurance. Testing the system should occur across the developmental and engineering life-cycle. Statistical methods such as design of experiments, sequential-adaptive test design techniques (e.g., response surface methods, optimal learning), and stratified random sampling for algorithms training, testing, should be employed to provide a systematic approach to scaling information collection with system development and ensuring adequate information exists across the various views on operating envelopes exists to support the assurance assessment. [2]

Execute & Analyze. The execution phase of testing should also include ongoing analyses of the data, assessments of feasibility, and lessons learned. Analyses should consider knowledge gaps based on planning strategies and dedicated time for reworking the test program based on information obtained during each phase of testing. [2]

Monitor & Improve. Our research recommends a continuous process of T&E for autonomous intelligent agents that shows that this data collection process is followed throughout development, manufacturing, deployment, and operations. In operations, they emphasize the need for both ongoing monitoring of the systems capabilities, independent assessment when pre-defined triggering events such as a mission objective change or an adversarial attack. Additionally, software defined functionality in general provides the opportunity to improve capabilities without a full redesign of the system. AI algorithms can be improved by introducing new more relevant data for the current tasking. Methods such as transfer learning and targeted fine-tuning provide the opportunity for continual improvement to fielded algorithms. However, robustness is always a consideration that must be considered when making algorithm improvements. The monitor phase is essential to consider from the early phases of system design. Monitoring is an easier task when the system is designed to output relevant measures and metrics automatically. [2]

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.

No Comment

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.

No Comment