

STANDARDIZATION EMPOWERING AI-ENABLED SYSTEMS IN HEALTHCARE

Report publication date: March 2021



WORKSHOP
REPORT



ANSI

The ANSI logo is located in the bottom left corner. It features the word "ANSI" in a bold, white, sans-serif font, centered within a light blue circle. The circle is partially overlaid by a white swoosh that curves around the top and left sides of the text.



TABLE OF CONTENTS

Introduction	3
Background for the Workshop	4
Workshop Objectives	5
Workshop Sessions	8
Overview of Artificial Intelligence Standards Landscape	8
Panel Discussion: Data (and All Its Demons)	8
Panel Discussion: Transparency AND Explainability	10
Panel Discussion: Governance and Risk Management	13
Recommendations	15
Appendix A: Participation	19
Appendix B: SDOs Referenced	20
Appendix C: Survey Questions	21
Appendix D: Workshop Agenda	22



INTRODUCTION

The landscape for health information and systems standardization and regulation is complex and very dynamic. This report highlights the primary stakeholders, resources, and themes in this subject area, the key challenges and issues that exist, and further areas for exploration related to the use of standardization to empower AI-enabled systems in healthcare.

EXECUTIVE SUMMARY

The importance and growth of autonomous and artificial intelligence (AI) systems has led to a strong consensus that ANSI should identify needs for the coordination of standardization for AI in healthcare, using a collaborative community approach. Current events and an explosion of AI in healthcare led to an initial broad stakeholder survey, a virtual workshop, and this report.

Survey responses identified wide-ranging uses of AI in consumer wellness, medical care, hospital operations, analytics, research, financing, cybersecurity, and administration. In many organizations and applications, AI is pervasive, oftentimes with apparent limited consumer awareness. Because AI is embedded in other technologies, AI powers and supports many things, and can help to enhance quality, manage complexity, and lower costs—high-level drivers in all healthcare sectors.

Much of AI standards development is sector agnostic, but healthcare has specific standards-related needs. Standard Development Organizations (SDOs) have an opportunity to be more nimble with AI technologies through partnerships with industry leadership, related regulatory agencies, and stakeholders. These communities strongly support ANSI principles and open consensus standards development processes.

Healthcare has unique considerations including sensitive health data, patient safety, and potential for disparate impact of AI. More standardization is needed to promote safety, effectiveness, and trustworthiness of healthcare-based AI and to support informed regulatory measures. Some states are enacting legislation focused on privacy, or automated decisions, that could set *de facto* national standards affecting healthcare AI.

Survey responses and workshop discussions recommended and encouraged some foundational principles and potential next steps that could be supported, coordinated, and developed by collaborations of ANSI with SDOs, NIST, and other government agencies. ANSI presently is working with stakeholders to develop plans for a future workshop to explore these needs, principles, and next steps for coordination of healthcare AI standardization.



BACKGROUND FOR THE WORKSHOP

To create a common understanding, identify and fill gaps in related work, and avoid unnecessary collisions or overlaps in standards for healthcare, ANSI agreed to sponsor a workshop seeking to identify challenges, gaps, opportunities, and collaborations, and to inform the development of regulatory frameworks. Particular focus was to be placed on the role of autonomous and intelligent systems (AI) in medical software, healthcare software, and consumer systems that interact with the professional healthcare sector. The goal of this effort was to develop recommendations for the coordination of standardization and governance so AI-enabled systems can meet expectations of safety, quality, responsibility, and risks.

A private/public planning committee first met in late February 2020 to explore this landscape with with co-conveners from Kaiser-Permanente (also a member of the ANSI Board) and NIST. During that meeting, planning began for a survey and an in-person workshop for autumn 2020. In March of 2020, the group discussed the change in situation due to COVID-19 and agreed that a face-to-face workshop would be unlikely in 2020 but moved ahead with planning to conduct a standards landscape scan, a leadership survey, and a virtual workshop to be held on September 14. The committee wanted to identify some of the characteristics and measures that would lead to trustworthy systems such as data quality, attributes of trustworthiness, and properties of data in the AI-enabled systems that would meet the needs of healthcare consumers, healthcare providers, and healthcare entities across the spectrum, as well as AI-enabled system developers, integrators, and other users.

A leadership survey was conducted in July 2020. It sought feedback from interested stakeholders on standardization issues, solutions, efforts, and the role of standardization in the governance and regulation of AI-enabled systems in healthcare. The survey results were used to inform the discussion at the September 14 virtual workshop and contribute to the themes and recommendations below. (Appendix C: Survey Questions)





WORKSHOP OBJECTIVES

The workshop's objectives were to identify challenges, gaps, opportunities, and collaborations, and to inform the development of regulatory frameworks. The goal was to have recommendations to improve standards coordination and governance seeking to assure stakeholders that AI-enabled systems can meet expectations of safety, quality, responsibility, and risks. The workshop considered the broadest scope of all the different types of AI—not just machine learning systems, but deterministic systems operating in a more constrained environment (both narrow and broad). Panelists were invited to discuss the quality of data and AI algorithms in software independent of devices as well as in medical devices. The themes that recurred throughout the multiple panels were data, trust, transparency, governance frameworks, and risk management. Based on survey comments, the workshop considered the full range of uses in healthcare and wellness, including:

- » Cybersecurity
- » Healthcare operations
- » Supply chain optimization
- » Workforce-related uses
- » Insurance underwriting
- » Revenue cycle management
- » Medical care
- » Clinical decision support
- » Radiology and other areas

While the workshop did not cover all of these topics, a subset of topics was addressed to meet the one-day workshop timeframe. The topics not addressed in the workshop should be further explored. The workshop objectives took a broad view of AI to identify other areas that would be helpful in identifying all the topics necessary to have trustworthy AI in healthcare. It was agreed that a common understanding based on standardized definitions of quality, explainability, transparency, privacy, and security would allow better management of the risk of these systems.

EMERGING THEMES

The following themes emerged from both the stakeholder survey and the September 14th Workshop:

Data

- » Data quality is critical to ensuring safe and effective medical AI systems that promote better patient outcomes.
- » Data in medical AI applications may be training, operational, or learning functions. Each of these functions may have specific needs, but semantic coherence and standardization of formats and normalizations will be required to ensure AI function.

- » Current challenges with data semantics are not solved with AI but, rather, may be magnified in AI applications.
- » In healthcare, there is machine learning for clinical decision support reporting, operational, financial, and population health uses. In all of these areas, data preparation is critically important in model development. Standards and measures will be needed to guide appropriate model development.
- » There are problems when operational data and biases are materially different from data that the system is trained on.
- » Scale really matters, both computationally and operationally in terms of the impact of both changes to data and perturbations in data.
- » Standards to provide easier data aggregation and preparation can leverage existing standards, such as HL7 FHIR application programming interfaces (APIs). Standards to enable better integration of models and consistency workflow can also leverage existing standards, such as HL7 standards.
- » Semantic interoperability and semantic standardization is a critically important issue for data normalization for model building and model use, and to be able account for error conditions. As a practical example, a lot of healthcare comes from using natural language processing (NLP) for chart abstractions. Do these data convey the original meaning adequately, and does the NLP introduce semantic model biases? What is the effect of the NLP models?
- » Future direction will focus on neuro-algorithmic hybrid models and quantum parallel models.

Trust

- » Trust is a complicated topic that can be based on the risk of the application and the human, technical, and regulatory factors.
- » The level of trust required will significantly vary based on the application by which the AI solution is being used.
- » Trustworthy models require trusted data that has appropriate provenance. This can include addressing topics of privacy, security, bias, and access to the data. There is a need to balance the technical applications of models and the human need to understand trusted data as a factor.
- » Human trust is going to be a key issue to the expansion of the use of AI in healthcare. There are many examples of lack of trust and even fear of AI in the public mind. A reoccurring theme is the central role of AI to automate administrative, repetitive tasks in healthcare. A higher level of attention to the trust dimensions of explainability and transparency is needed when AI provides care management or health management recommendations
- » How can one trust the results and output of AI-enabled systems if they don't understand them? How can the use of AI be governed if it is not understood? There is a need to consider an approach to explainability or transparency that allows for humans to understand what the AI is doing. This will be a critical issue when considering where there is human intervention with the AI system.

Transparency

- » Transparency is key to understanding data and algorithms; this understanding is required to build trust in them.

- » Transparency is needed in regulatory data and in algorithms for the “who” and “how.” Transparency is also needed in independent review. An especially challenging aspect of testing and implementing AI-enabled systems is understanding the model itself. System complexity and the dynamic nature of machine learning make it challenging to fully explain how the AI-enabled system processes data and produces particular results.
- » Develop a 5-year plan and a 10-year plan for transparency, which includes by whom and how models were developed, for independent review of recommendations or model outputs.

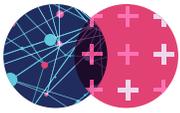
Governance Frameworks

- » Organizations benefit from using a governance framework¹ to guide their decisions in how to use and rely on AI in various aspects of their organizations. Governance frameworks for AI management are in early stage use. There is a desire to have best practices, voluntary guiding principles, and voluntary consensus standards as parts of a governance framework.
- » Policies that may be governed by this framework could address the culture of safety, such as how voluntary reporting of errors and safety risks in healthcare are applied in the context of AI. Policies on disclosure of design failures, error reporting, and disclosure of biases when discovered are needed.
- » It will be important to have consistent measures and definitions of fairness, disparate impact, and a consistent understanding of the characteristics of acceptable and unacceptable decision criteria across a wide of variety of different specialty cases.
- » A governance framework should address consistent adoption of standards and measures, including privacy policies enabling individuals to know how their data will be used.

Risk Management

- » What risk management approaches are suitable for managing the risk of AI use? What new approaches are needed? There was a strong discussion on adapting the existing ways of managing risk with new risks—we don't need to start over.
- » It is necessary to consider the real world environment's impact on the use of AI models, including user interactions with the systems and the physical and logical environment in which they are used.
- » There is a repeated theme of striking the right balance, i.e. balancing costs with benefits, being thorough but not overly complex in review.
- » Stakeholder education is key: stakeholders must understand how things work (and limitations to that work).
- » All the right stakeholders must be involved.
- » Clear guidance on privacy and consent language is needed.

¹ Governance frameworks can be defined as: “The system by which entities are directed and controlled. It is concerned with structure and processes for decision making, accountability, control and behaviour at the top of an entity or effort. Governance influences how objectives are set and achieved, how risk is monitored and addressed and how performance is optimised.” Governance is a system and process, not a single activity and therefore successful implementation of a good governance strategy requires a systematic approach that incorporates strategic planning, risk management, and performance management. Like culture, it is a core component of the unique characteristics of a successful organisation or effort.” Source: [governancetoday.com](https://www.governancetoday.com)



WORKSHOP SESSIONS



OVERVIEW OF ARTIFICIAL INTELLIGENCE STANDARDS LANDSCAPE

Presenter: Pat Baird, Sr. Regulatory Specialist—Head of Global Software Standards, Philips

Mr. Baird presented on the collective knowledge and experience in the AI standards landscape, and reported on “pre-standards” activities that could help identify sub-topics of interest and potential starting points for future standards. Standards need to maximize the advantages of AI in advancing patient health by identifying how to provide a reasonable level of confidence in the performance of AI-enabled systems in a way that minimizes risks to product quality and patient safety. Mr. Baird’s overview included efforts from AAMI/BSI, Consumer Technology Association (CTA), IEC, IEEE, ISO/IEC JTC1, ISO/TC 215, and UL. (Appendix B: Standards Development Organizations Referenced)



PANEL DISCUSSION: DATA (AND ALL ITS DEMONS)

Moderator: Joe Lewelling, Senior Advisor on Content and Strategy, AAMI

Ittai Dayan, Executive Director of Strategy and Operations, MGH & BWH Center for Clinical Data Science, Partners Healthcare

James M. Dzierzanowski, Executive Director, Innovation and Strategic Advisory Services, Kaiser Permanente

Wayne Kubick, CTO, HL7

Dr. Tanju Gupta, VP of Cerner Intelligence, Cerner Corp.

The panel discussed healthcare data challenges, agreeing that sufficient accuracy and quality of data are especially necessary for AI-enabled medical technology—both for basic rules-driven systems and for systems that employ more advanced machine learning. Furthermore, data in AI systems may have different functions—the data may be used to initially train the system (*training data*), it may be used to determine a course of action for an individual patient (*operational data*), or the data may be used by the system to modify or refine treatment to achieve better outcomes with future patients (*learning data*). The requirements for data can vary across this *training–operational–learning* continuum depending on the data’s intended function and the context of its use.

POINTS, CHALLENGES AND FUTURE QUESTIONS ON DATA

The panelists discussed unique challenges with respect to data utilized by AI-enabled machine learning systems.

Data Quality Challenges

- » When an AI system can propose diagnoses or base treatment decisions on data (perhaps without direct clinical review), the criticality of data preparation and quality is elevated.
- » Proper data selection and recording of its provenance is essential to ensuring data trustworthiness as well as its adequacy and scalability for use across the larger patient population.
- » All data transformations or normalizations must be understood and standardized to ensure that training, operational, and learning data are coherent and treated consistently (appropriately) by the medical AI system.
- » Data production, acquisition, and use involve both humans and technology, and ensuring the semantic coherence of that data—that it has the same logical and contextual meaning to both the people and the AI involved—is crucial to maximize the effectiveness and minimize the risks of AI systems.
- » Error conditions and error codes must also be consistent so that they can be properly handled and assessed by both the systems and by users.

Data Analysis Challenges

- » AI potentially can perform complex predictive analyses that will drive unprecedented improvement in healthcare delivery. To achieve these improvements, the data sets utilized in training and learning must be both adequate and representative of the patient populations.
- » The ability of an AI system to recognize when data sets do not meet required conditions due to selection bias or other confounding factors must be aligned with its ability to undertake analyses and act autonomously.
- » Data sets may be biased in terms of demographics like race, gender, ethnicity, or socioeconomic status. This can be a result of inadequate collection practices or of self-selection by members of the target populations.
- » Similarly, healthcare inevitably collects more data on sick people than on healthy people. Data will require curation externally or by the AI system to ensure biased data sets do not skew data analyses.
- » Data drift—the tendency of data sets to evolve over time due to, for example, changes in disease presentation or changes in the patient population—can be problematic if an AI system is not capable of recognizing and accommodating such drift.
- » Conversely, properly modelled AI systems potentially could be better positioned to recognize and deal with problems like non-representative data sets and data drift than “unintelligent” static rules-driven systems.

OPPORTUNITIES FOR STANDARDIZATION

The panel discussed the current state of data and data set standardization. The AI industry is moving at a highly innovative pace, but its use in healthcare delivery is still somewhat emergent. Standard data frameworks, algorithm models, and terminologies exist, but must be adapted to align with AI requirements. Measurable characteristics of data sets of specific relevance to AI are a standardization opportunity.

The panelists noted that healthcare data entry is collected from multiple points in the system. The shifting of data entry to the machine reduces burden on human entry and the potential for human error, but it is necessary for the data entry itself to be standardized. Data visibility through APIs may also help—over time, data sets can be more consistent with improvements.

Critical to the future success of medical AI is not only additional standardization of AI model characteristics but also the development of standardized model language and requirements for data selection, collection, curation, and quality necessary to support each of those models.

Medical AI may have unique data needs, but standardization in this area must take place within context of the larger healthcare delivery enterprise and with consideration of real-world needs and limitations.



PANEL DISCUSSION: TRANSPARENCY AND EXPLAINABILITY

Moderator: Heike Sichtig, Ph.D., SME & Team Leader, Digital Health, FDA

Dr. Andrea M. Matwyshyn, Professor, FDA

Dr. Jesse M. Ehrenfeld, Chair, AMA

Suzanna Hoppszallern, Senior Editor Data and Research, AHA

The panel, moderated by Dr. Heike Sichtig, discussed the “intent” (goals/performance) of AI-enabled systems compared to their “output” (accuracy/validation), and the potential regulatory and non-regulatory impact.

POINTS, CHALLENGES AND FUTURE QUESTIONS ON TRANSPARENCY AND EXPLAINABILITY:

One key principle noted was the need for collaboration and coordination among stakeholders.

Concerns in development of AI-enabled systems include:

- » The addition of technology capabilities just because they are available but that do not necessarily add value or do not benefit the end-user
- » Builder bias that results in quick deployments where security is left out of the analysis and where the players at the table are impeded by their knowledge limitations

- » Users' "over-trust" of technologies and input from devices
- » Security maturity models are not currently incorporated into aspects of threat modeling for AI.
- » AI should be using cybersecurity frameworks and standards already developed (i.e. NIST, ISO standards).

Clinician challenges:

- » Physicians need to understand the AI methods or systems in order to trust the algorithmic predictions. They need explainable algorithms.
- » How do we know we are not incorporating additional bias by the application of an AI system?
- » In some cases, it is the responsibility of the physician from an ethical standpoint to obtain valid, informed consent from the patient on the use of a product or tool. The physician needs to provide relevant information to the patient on the benefits, risks, and alternatives. Explaining the role of the AI solution to the patient, including its risks and benefits, should be managed appropriately.
- » The novelty and sophistication of AI tools place additional demands on an informed consent process. The presentation of AI information can be complicated, which can result in overconfidence.
- » The physician has to be sufficiently knowledgeable to understand how the AI system is working. This is rendered difficult with systems that lack transparency or explainability.

For transparency, there is a need to involve clinicians in reviewing processes, methodologies, curation, integration, and ethical decisions of AI systems and tools. Clinicians must be involved in the development and design of the AI clinical solution from start to finish. Transparency in the AI algorithm to the healthcare work force is also necessary in order to understand what went into it and why.

For explainability, there is a need to use validated frameworks and learn from successes and failures. The data must show that the AI technology solves the clinical problem it was designed for in real-world clinical settings.

Given the current level of autonomy in healthcare applications, some people use the phrase "augmented intelligence" or "assisted intelligence" instead of "artificial intelligence," emphasizing that software is supporting—not replacing—caregivers.

AHA shared their perspective that AI is part of the healthcare transformation (just as it has transformed other segments), and they want to see AI solving problems of affordability, access, outcomes, and experience. There are many examples where hospitals and healthcare system are using AI in operations and administration. AI will not take over decision-making—it will augment care and become part of the healthcare team as a powerful tool in the toolbox.

AHA has concern about the administrative burdens on clinicians and physicians and nurses. However, as an example, voice recognition is a potential opportunity to reduce burden, as well as automating repetitive administrative tasks such as billing. It was noted that individuals are taking responsibility for their own health as they are feeling the pressures from the digitized environment.

Additionally, hospitals are looking at ways to improve care early on—e.g., clinical decision support tools for kidney disease.

Further areas of concern were also discussed. Some physicians are concerned that AI systems must promise a better outcome and experience, and provide help in managing excessive workloads, in order to benefit care.

AI should be adequately explainable and physicians must be involved in AI systems. Some patients are skeptical about using AI, but are becoming more comfortable. They need assurance that AI is trustworthy. AI was used to triage many patients during the pandemic and was very helpful. A future need is to focus on personalizing and individualizing the healthcare experience. Another element is ethical considerations, including considerable concerns about unintended/unknown bias or disparate impact. There should be awareness of bias based on race, etc. Also, patient privacy is paramount, and it affects how AI is developed and tested. Transparency in the algorithm is necessary.

Parameters for transparency include performance characteristics, population in which the products were tested and those in which it has not been tested, and explainability. (Device characteristics should be explained in language that the user can understand; this may require a new approach to labeling.)

OPPORTUNITIES FOR STANDARDIZATION AND FURTHER CONSIDERATIONS:

- » The panel discussed standardization needs such as default settings fields and the disclosure around points of failure. When simulation data is created, there is an opportunity for transparency.
- » Are standards needed for de-identified repurposed data? Patients need to know when external bodies process/sell their information and what benefits may come from that—for example, improving data for more informed decisions.
- » Design choices versus patient safety should never be a choice (e.g., sensors in bodies collecting data).
- » How can clinicians communicate to manufacturers? How are they informed when there is an update to systems algorithms?



PANEL DISCUSSION: GOVERNANCE AND RISK MANAGEMENT

Moderator: **Dr. John Halamka**, President, Mayo Clinic Platform

Marilyn Zigmund-Luke, VP, AHIP

Amy Mahn, International Policy Specialist, NIST

Pat Baird, Sr. Regulatory Specialist, Head of Global Software Standards, Philips

John Dickerson, Assistant Professor in the Department of Computer Science, University of Maryland

Dr. Halamka opened the session with some comments describing how his organization, Mayo Clinic Platform, as a large healthcare ecosystem, curated data to build AI algorithms and de-identify data for privacy. He illustrated that when patients were informed how data was de-identified transparently, the patients felt their privacy was safe. When the data is de-identified, algorithm developers are invited in to access the closed container data via API access but cannot infiltrate the data and it never leaves Mayo's control. The Mayo Clinic Platform has a multiple-layer governance structure in place, including security infrastructure, controlled data use, and audits that supports the ethical uses of data. Their data is privacy preserving, reasonable, and governed. Mayo has worldwide discussions to develop and revise their consent language to verify its legality and cultural appropriateness worldwide.

The panel discussed a need for governance of AI (at the organizational level) in healthcare with regard to privacy, risk management, and efficacy. Governance of the systems needs to be determined to manage processes. Panelists outlined considerations to determine if risk management frameworks are necessary.

Governance of AI-enabled systems should respect agreed principles, and in areas such as:

- » **Privacy:** AI systems should respect individual privacy and enable safe, effective healthcare that requires the use of personal data.
- » **Accountability:** AI systems should be auditable, and their impacts should be appropriately identified and distributed, with mutually agreed remedies as needed.
- » **Safety:** AI systems should be safe, performing only as intended.
- » **Security:** AI systems should be secure from compromise by unauthorized parties.
- » **Transparency and Explainability:** Design and implementation of AI systems should enable transparent operations that provide information about where, when, how, and for what purposes they are being used.
- » **Fairness and Non-discrimination:** AI systems should have mechanisms to address concerns about bias in their data or results, for fairness and inclusivity.
- » **Human Control of Technology:** Final decisions should remain subject to human review.
- » **Professional Responsibility:** Individuals who develop or deploy AI systems should involve appropriate stakeholders and consider long-term effects or unintended consequences.

- » **Effectiveness:** The degree to which an AI-enabled system achieves its intended outcome should be measured and managed.

POINTS, CHALLENGES AND FUTURE QUESTIONS ON GOVERNANCE AND RISK MANAGEMENT:

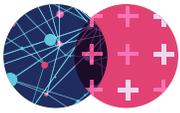
Panelists agreed that consent language is necessary and regulations shouldn't make processes overly complex or hinder the AI-enabled systems. Consider balancing costs with benefits, but there must be an option for human review for oversight. Over-trust can be caused when there is not a familiarity with AI, and this needs to be balanced with under-trust, when one is suspicious of AI-enabled system results. Striking a balance here is supported by explained governance of the systems.

Band its disparate impact factors into AI-enabled systems. Discrimination can be system learned, so mitigation is important. Fairness needs to be defined and encoded into metrics. There was agreement on constant monitoring and continuously refining metrics.

The criteria and risk categorization framework developed by the International Medical Device Regulators Forum and adopted by the FDA for determining software that should receive medical device oversight sets a standard for the extent to which transparency and explainability are required, and under what conditions. While this framework is intended for AI-enabled systems used by clinicians and patients in the medical decision-making process, it can serve as a model for developing criteria and assessing risk of AI models used in other healthcare domains.

OPPORTUNITIES FOR STANDARDIZATION:

The panelists agreed that public/private dialogue is necessary and all stakeholders need to be involved. Consumers also need education to determine their risk assessment. Guiding principles in organizational governance should be developed. Organizations need a risk-based approach as well as measures to define fairness and disparate treatment. The first step is determining if a group is adversely impacted by a defining characteristic, and if the decision was made based on membership in that group. When that is determined, there should be framework to change that impact. Organizations should have fairness standards and a consistent review of those in the modelling. Ethics and bias are large concerns and most of the data sets are not large enough to test for age, gender, or race bias. It is necessary to know if an algorithm will introduce economic disparities or a bias against certain groups.



RECOMMENDATIONS

This section summarizes areas for further exploration in the context of the themes identified.

DATA

- » There is a need to develop standards for characteristics of data sets in AI to be used in assessments, measurements, and practices including roles, activities, and responsibilities.
- » Data quality is of high importance; although specific applications will have specific data quality requirements, a general approach to ensuring quality data should be developed and should contain consensus-based standards and practices.
- » Data quality frameworks should account for characteristics and requirements for various data types, training data, learning data, and operational data (such as handling, legal, security, retention, access, provenance, demographics, and authorization, among other characteristics).
- » Data lifecycle frameworks in organizational policies should include components such as initial acquisition, data cleansing, transformation, quality assurance, tagging, storage and retention requirements. These may need to be domain-specific standards for certain types of data (e.g. characteristics of good genomic data may be different than good vital signs data). General principles can be established in a generic framework.
- » Data quality considerations should be addressed, such as bad data coming in (bad sensors), how it can adversely affect data, and how it can be cleaned.
- » Guidelines for good synthetic data practices should be developed. Synthetic data is sometimes used as test data or research data to help protect patient privacy. It is also sometimes used in situations where additional data points are required (e.g. analysis of a rare disease).

TRUST

- » Trust is a complicated topic that can be based on the risk of the application along with the human, technical, and regulatory factors. Trust frameworks should recognize that the level of trust required can vary significantly based on the application by which the AI solution is being used. For example, data being used in a wellness application may be perceived differently than a hospital-controlled device monitoring application.
- » Trustworthy models require trusted data that has appropriate provenance. Trusted models should address topics of privacy, security, bias, and access to the data.
- » There is a human need to understand trusted data through transparency that should be balanced with the needs of technical applications.
- » A higher level of attention should be paid to explainability and transparency to address the lack of trust of AI in the public mind, especially where AI provides care management or health management recommendations.
- » The relationship of human intervention to trust should be further explored.

TRANSPARENCY

- » Transparency can increase public trust and confidence in AI applications.
- » Standards and policies for transparency can inform regulatory decisions, individual awareness, purchasing decisions, and organizational policies.
- » The evolving nature of AI solutions necessitates a multi-year plan for transparency standards.
- » Standards related to transparency should consider the relationship between transparency, trust, privacy considerations, and commercial confidentiality, by whom and how models are developed, as well as requirements for independent review of model outputs.

GOVERNANCE FRAMEWORKS

- » The roles and relationships of AI governance frameworks and AI risk management frameworks should be well defined with agreed concepts and terms, integrated and coordinated with other frameworks.
- » The community should leverage existing governance and risk management frameworks that are suitable for AI. A governance framework should rely on consensus-based measures and definitions of fairness, disparate impact, and a consistent understanding of the characteristics of acceptable and unacceptable decision criteria across a wide of variety of different specialty cases, and should address the culture of safety with voluntary reporting of errors and safety risks, disclosure of design failures, error reporting, and biases in the context of AI.
- » A governance framework should help organizations meet policy and regulatory requirements; however, these requirements should not be part of the framework.
- » Voluntary guiding principles and a governance framework are important and should inform regulations. They should be adoptable by organizations in both the private and public sectors, addressing common needs.
- » Consistent adoption of frameworks for standards and measures should address provisions enabling individuals to know how their data will be used.

RISK MANAGEMENT

- » The roles and relationships of AI risk management and AI governance frameworks should be well defined with agreed concepts and terms, integrated and coordinated with other frameworks. Many of the governance recommendations also apply to risk management.
- » Risk management frameworks should be clearly defined and should address the needs of patient risks as well as organization risks.
- » A risk framework should help to determine which risks are acceptable and which risks present the possibility of unacceptable harm, or harm that has expected costs greater than expected benefits.
- » Risk management should address issues including: communication and overall stakeholder education, a taxonomy of AI failure modes, risk controls, software defects, data quality, bias, privacy, cybersecurity, failures of synthetic datasets, trust, levels of autonomy, and failures related to adaptive systems.

- » Reliance on human intervention and operational control should be considered in risk management frameworks.
- » Exploration is needed regarding assumptions about individuals based on incomplete data, related to when individuals need to be aware of information gathered about them and its uses.

NEXT STEPS/FUTURE WORKSHOPS

Plans are in place for a hybrid (in person/virtual) workshop in late 2021 (date TBD) to take a deeper dive to explore challenges and opportunities in standards development.

Visit [this page](#) for updates.

APPENDICES





APPENDIX A: PARTICIPATION

The following organizations participated in this workshop meeting:

<i>98point6</i>	<i>Draegerwerk AG & Co. KGaA</i>	<i>Oracle</i>
<i>AAMI</i>	<i>Eli Lilly and Company</i>	<i>PALM Associates, Inc.</i>
<i>American Dental Association</i>	<i>FDA</i>	<i>Philips</i>
<i>America's Health Insurance Plans</i>	<i>FITSCRIPT</i>	<i>PJM Consulting LLC</i>
<i>American Hospital Association (AHA) Center for Health Innovation</i>	<i>GS1 US</i>	<i>Sechenov University</i>
<i>American Medical Association</i>	<i>Hillrom</i>	<i>Sullivan County Community Hospital</i>
<i>American National Standards Institute (ANSI)</i>	<i>HL7</i>	<i>The Joint Commission</i>
<i>Americans In Need</i>	<i>IEEE</i>	<i>Trusted Solutions Foundry</i>
<i>BAH German Medicines Manufacturers Association</i>	<i>InfoBeans Inc.</i>	<i>Tufts School of Medicine</i>
<i>BDC-HealthIT</i>	<i>Inteleos</i>	<i>Tufts University</i>
<i>CareMoat Inc.</i>	<i>Intertek Consulting Services</i>	<i>TÜV SÜD Americas</i>
<i>Cerner Corporation</i>	<i>Kaiser Permanente</i>	<i>UK Health Care</i>
<i>Consumer Technology Association</i>	<i>LCI</i>	<i>United Spinal Association</i>
<i>ConvergeHEALTH by Deloitte</i>	<i>Mayo Clinic Platform</i>	<i>University of Maryland</i>
<i>Department of Veterans Affairs</i>	<i>MGH & BWH Center for Clinical Data Science</i>	<i>VA/Book Zurman</i>
	<i>National Institute of Standards and Technology (NIST)</i>	<i>Workcred</i>



APPENDIX B: STANDARDS DEVELOPMENT ORGANIZATIONS REFERENCED

Association for the Advancement of Medical Instrumentation (AAMI)

Consumer Technology Association (CTA)

International Electrotechnical Commission (IEC)

Institute of Electrical and Electronics Engineers (IEEE)

International Organization for Standardization/IEC Joint Technical Committee 1, Information Technology, subcommittee 42, Artificial intelligence (ISO/IEC JTC1/SC 42)

ISO Technical Committee 215, Health informatics (ISO/TC 215)

National Institute of Standards and Technology (NIST)

Underwriters Laboratories (UL)



APPENDIX C: SURVEY QUESTIONS

1. How are you currently using Artificial Intelligence (AI)-enabled systems in healthcare?
2. What barriers do you see to adoption, implementation, or use of AI-enabled systems in healthcare?
3. What characteristics would be most conducive to ensuring that AI-enabled systems could be governed appropriately? For example:
 - » Parameters for transparency, or
 - » Definitions and characteristics of trustworthiness
4. What characteristics related to data quality are important for standardization of AI-enabled systems, and for what purposes? For example:
 - » What characteristics of data could detect or measure misclassification of data (either intentional or unintentional misclassification)?
 - » What characteristics of data could detect or measure otherwise unknown or unwanted bias in data sets used as inputs, or in data outputs of AI-enabled systems?
 - » What characteristics of data could detect or measure data poisoning, fraud, manipulation, perturbations, or adversarial inputs?
 - » How should control and provenance of data sets be managed in AI enabled systems?
5. What standardization could enable lifecycles of AI-enabled systems to be measured or managed appropriately? For example:
 - » What are the characteristics of algorithmic obsolescence?
 - » What lifecycle management methods or frameworks should be supported by standard definitions or measures for AI-enabled systems?
 - » What lifecycles standards are needed for AI-enabled systems that are different from other systems?
6. How can privacy risk be managed in a way that supports the validity of the data sets of AI-enabled systems while protecting individual privacy, and what is the relationship of data privacy to trustworthiness of AI-enabled systems?
7. How do safety considerations relate to standardization necessary for AI-enabled systems?
8. What would be most useful in regulatory risk management and regulatory reporting for AI-enabled systems?



APPENDIX D: WORKSHOP AGENDA

SEPTEMBER 14, 2020 | 1:00 PM—4:15 PM EASTERN

Standardization Empowering AI-enabled Systems in Healthcare Coordination Workshop

As the national coordinator for the U.S. private-sector system of voluntary standardization, the American National Standards Institute (ANSI) has convened this meeting to explore opportunities for progress through collaboration and standardization, identify challenges, barriers and gaps, and discuss steps to optimize regulatory frameworks.

The goal of this virtual workshop is to set the stage and momentum for an in person workshop in 2021, where recommendations for coordination of standardization and governance to meet expectations of safety, quality, responsibility, and risk can be developed.

TIME	DISCUSSION TOPIC AND SPEAKERS
1:00—1:10 pm	<p>Welcome and Opening Remarks</p> <p><i>Jamie Ferguson, VP, Health IT Strategy & Policy, Kaiser Permanente</i></p> <p><i>Lisa Carnahan, Assoc. Dir. for IT Standardization, Information Technology Laboratory, NIST</i></p>
1:10—1:25 pm	<p>Overview of Artificial Intelligence Standards landscape</p> <p><i>Pat Baird, Sr. Regulatory Specialist - Head of Global Software Standards, Philips</i></p> <p>Mr. Baird will provide an overview of the current standards landscape as well as projects under development.</p>
1:25—2:10 pm	<p>Panel: Data (and all its demons)</p> <p><i>Moderator: Joe Lewelling, Senior Advisor on Content & Strategy, AAMI</i></p> <p>The requirements and issues for operational data and training data are very much intertwined with respect to AI because of the ML-aspect. When and where they differentiate will be explored along with issues of measuring data quality, bias in inputs and outputs, data classification, and semantic standardization.</p> <p><i>Panelists:</i></p> <p><i>Ittai Dayan, Executive Director of Strategy and Operations, MGH & BWH Center for Clinical Data Science, Partners Healthcare</i></p> <p><i>James M Dzierzanowski, Executive Director, Innovation and Strategic Advisory Services, Kaiser Permanente</i></p> <p><i>Dr. Tanuj Gupta, VP of Cerner Intelligence, Cerner Corp.</i></p> <p><i>Wayne Kubick, CTO, HL7</i></p>

TIME	DISCUSSION TOPIC AND SPEAKERS
2:10–2:55 pm	<p>Panel: Transparency AND Explainability</p> <p><i>Moderator: Heike Sichtig, Ph.D., SME & Team Lead, Digital Health, FDA</i></p> <p>What is the “Intent” (goals/performance) compared to “Output” (accuracy/validation) and the Regulatory Impact?</p> <p><i>Panelists:</i></p> <p><i>Dr. Andrea M. Matwyshyn, Professor, FDA</i></p> <p><i>Dr. Jesse M. Ehrenfeld, Immediate Past Chair, AMA Board of Trustees</i></p> <p><i>Suzanna Hoppszallern, Senior Editor Data and Research, American Hospital Association</i></p>
2:55–3:10 pm	Virtual Break
3:10–3:55 pm	<p>Panel: Governance and Risk Management</p> <p><i>Moderator: Dr. John Halamka, President, Mayo Clinic Platform</i></p> <p>There is a need for governance of AI (organizational level) in healthcare with regard to: Privacy, risk management, and efficacy</p> <p>How to govern the systems for managing processes? Is a risk management framework needed?</p> <p><i>Panelists:</i></p> <p><i>Marilyn Zigmund-Luke, VP, AHIP</i></p> <p><i>Amy Mahn, International Policy Specialist, NIST</i></p> <p><i>Pat Baird, Sr. Regulatory Specialist - Head of Global Software Standards, Philips</i></p> <p><i>John Dickerson, Assistant Professor in the Department of Computer Science, University of Maryland</i></p>
3:55–4:15 pm	<p>Next Steps and Plans for 2021 Workshop</p> <p><i>Jamie Ferguson, VP, Health IT Strategy & Policy, Kaiser Permanente</i></p> <p><i>Lisa Carnahan, Assoc. Dir. for IT Standardization, Information Technology Laboratory, NIST</i></p>
4:15 pm	Adjourn

FOR MORE INFORMATION, CONTACT:

Michelle Maas-Deane

25 West 43rd Street, Fourth Floor

New York, NY 10036

T: 212.642.4884

E: mdeane@ansi.org

