Update on Global AI/ML Regulations and Standards

Brandon O'Shea (GE HealthCare)

AI SUMMIT

CINCINNATI, OH · NOVEMBER 14–16, 2023

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Inspiring Collaboration. Leading Innovation. Making a difference.

Global View | Emerging Laws Meeting Existing Frameworks

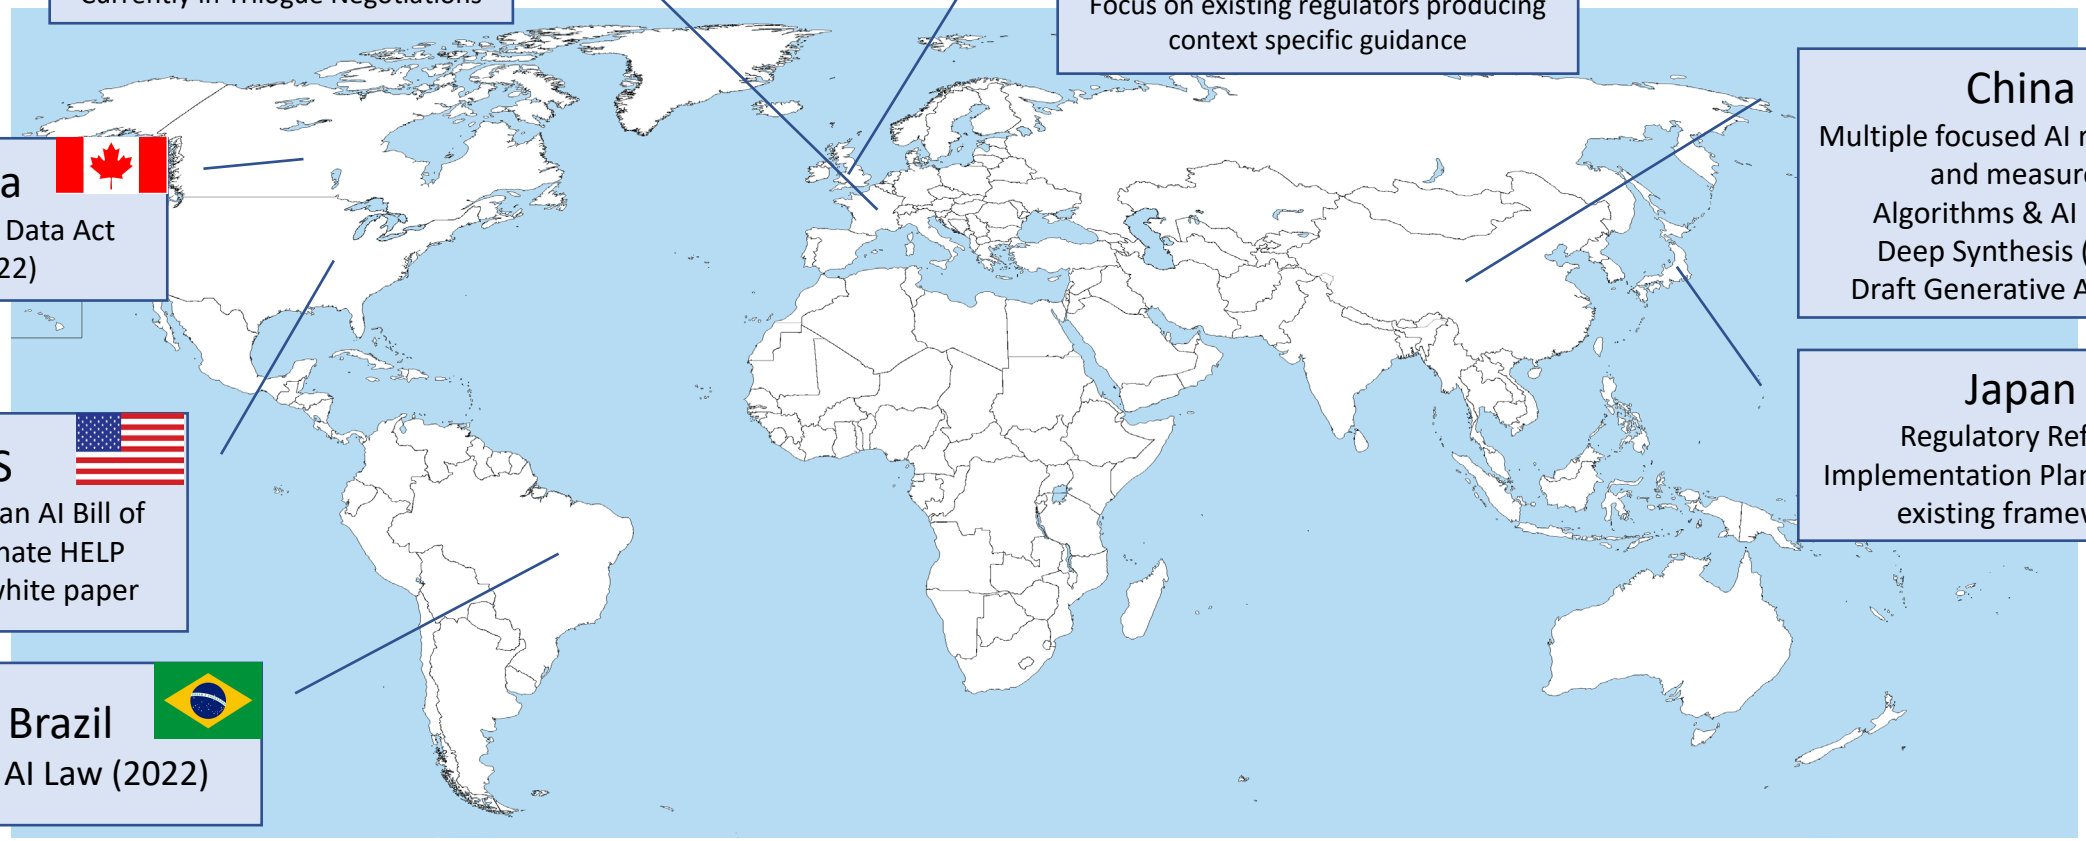

EU
(Draft) EU AI Act (AIA) (2021)
Currently in Trilogue Negotiations

UK
AI White Paper
Focus on existing regulators producing context specific guidance

Canada
Proposed AI and Data Act (AIDA) (2022)

China
Multiple focused AI regulations and measures
Algorithms & AI (2021)
Deep Synthesis (2022)
Draft Generative AI (2023)

US
Blueprint for an AI Bill of Rights & Senate HELP committee white paper

Japan
Regulatory Reform Implementation Plan (2021) in existing framework

Brazil
(Draft) AI Law (2022)

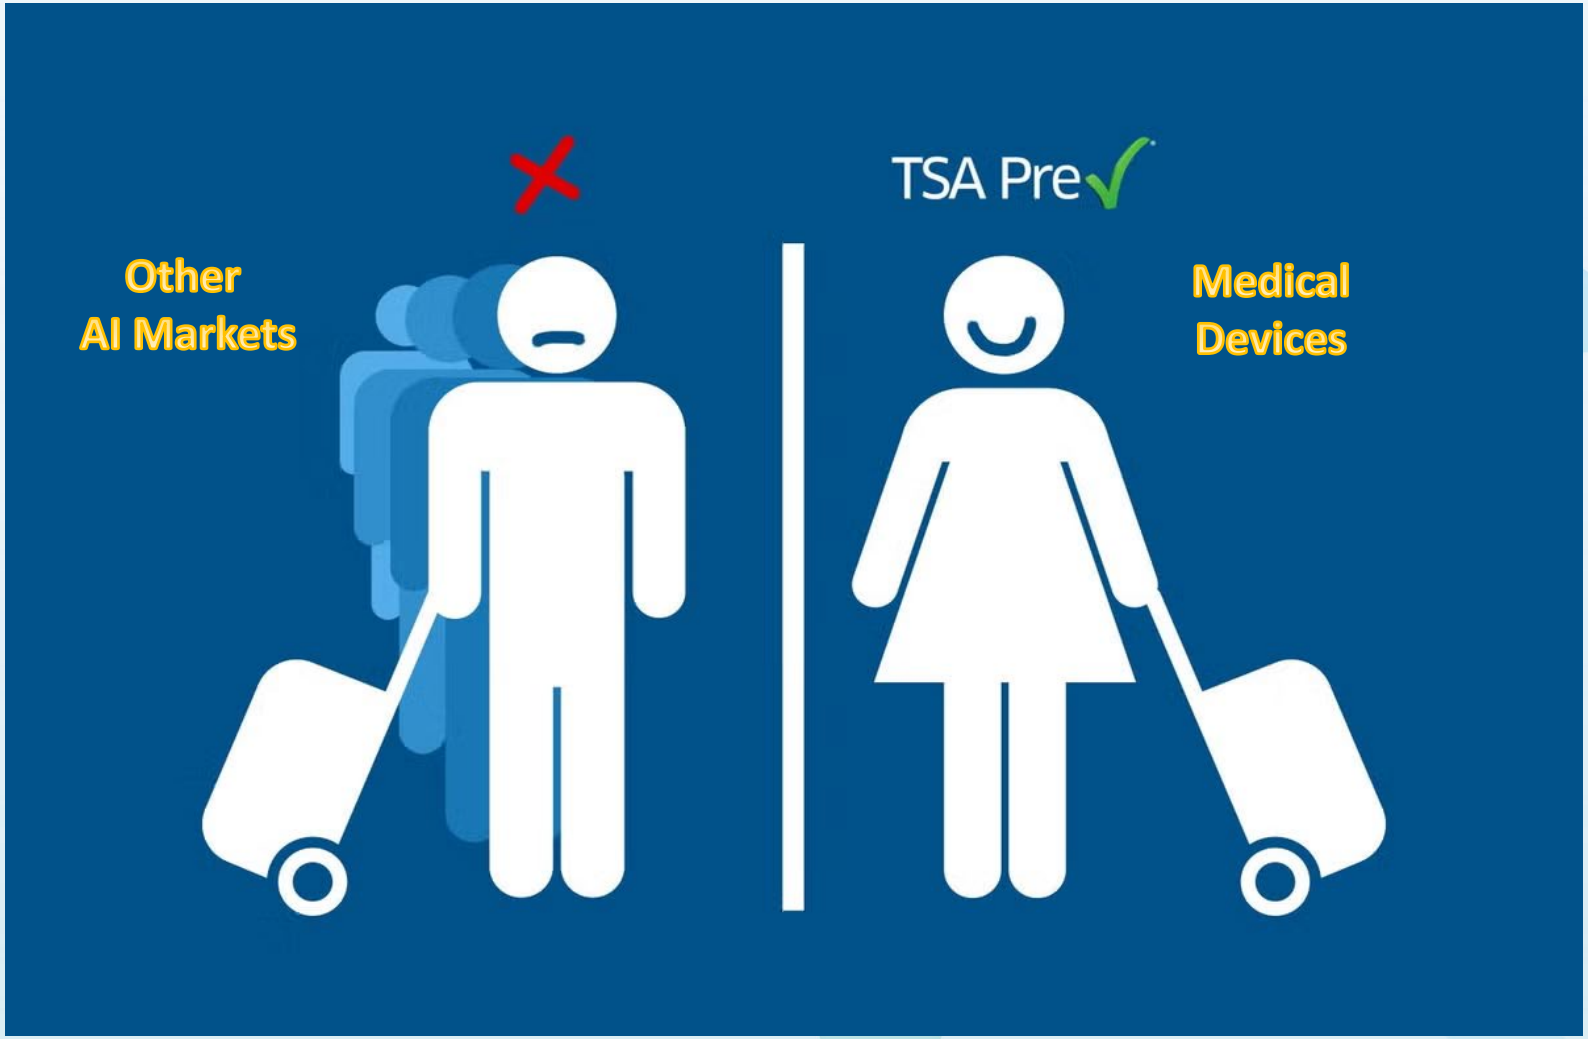
Other
AI Markets
TSA Pre
Medical
Devices

Framework Expansion | European Union

Existing MD framework seeking to show AI elements can fit in

Consumer Protection Driven Framework

Comprehensive (Horizontal) Risk Based Approach to AI

- Stating the framework first with the vertical pillars for implementation to continue to be developed and sorted out

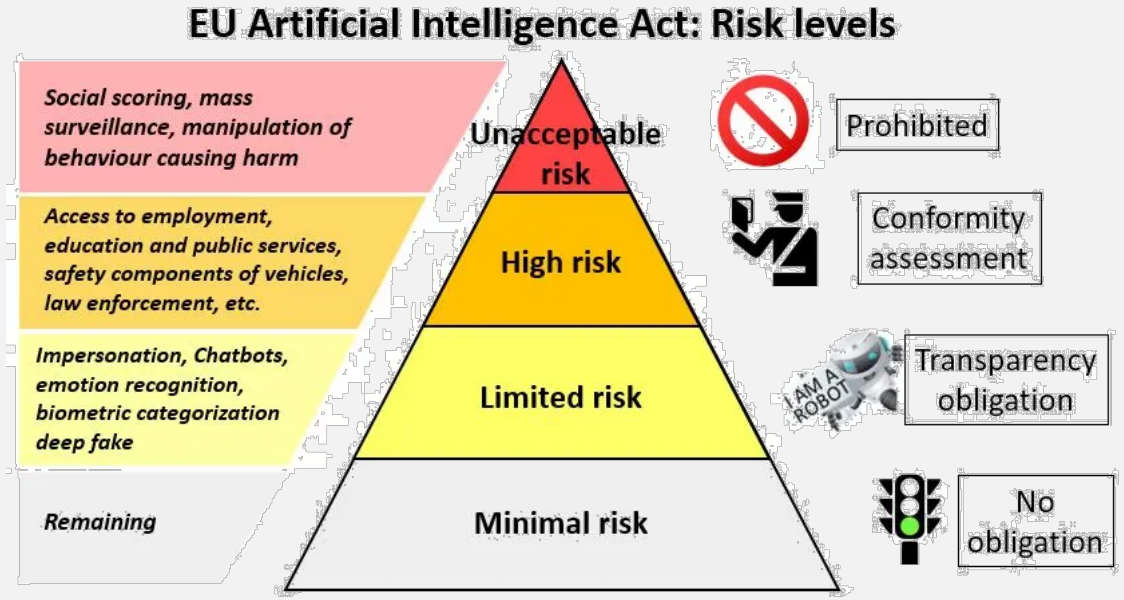

Existing Practices Adopting AI Elements

Conformity Assessment for High-Risk AI Systems:

- Case made to utilize EU sectoral law (e.g. MDR) with 3rd party conformity assessment reviews

Elements which build on existing framework:

- Risk Management, QMS, Essential Requirements
- Combined Technical Documentation & Declarations of Conformity (MDR, AIA, others)

Newer elements:

- Article 12 – Record Keeping (automatic recording of events ('logs'))
- Article 13 – Transparency to Users

AI SUMMIT
CINCINNATI, OH · NOVEMBER 14–16, 2023

Framework Review | United States

Taking the existing framework forward with clarity and innovation

United States (FDA)

Approach remains focused on the existing focus on intended use, while utilizing guidance documents to inject clarity and innovation

Following Through on Commitments

- FDA and the Digital CoE are continuing to develop the guidance and pathways for innovative AI/ML
  - 2022 CDS Guidance
  - 2023 *Draft* PCCP Guidance
    - QVD – Rad. ML Based Quant. Imaging SW with a PCCP (including adaptive vs QIH non-adaptive)

The A-List Priorities (FY2024)

- Final AI/ML PCCP Guidance
- *Draft* AI/ML Software Lifecycle Management Consideration (premarket)
- *Draft* PCCP for Medical Devices (general)

United States (HELP Committee)

Reviewing the Framework for Improvements

FDA generally well-suited to adapt to the use of AI

- Concern whether the framework needs review with regard to unlocked (adaptive/evolving) AI applications
- Foundational questions regarding transparency, development, ongoing effectiveness and liability boundaries
- Standards needed to demonstrate clinical validity of AI

Focus on FDA review speed and efficiency

- Encouraged to be early adopters in use of AI-powered tools in the process

Framework Buildup | China (NMPA)

Continuing to quickly build targeted AI regulations

NMPA

Building up from the foundation (2021 Order 739)

Remains focused on intended use driving approach but approaching it through categorization

Recent Guidelines

July 2023 four (4) new guidelines for AI software:

- Ultrasound Imaging AI Software (Process Optimization)
- Performance Evaluation on AI Analysis SW for Pathological Images
- Clinical Evaluation on AI Analysis SW for Pathological Images
- Performance Evaluation on AI Analysis SW for Blood Disease Flow Cytometry

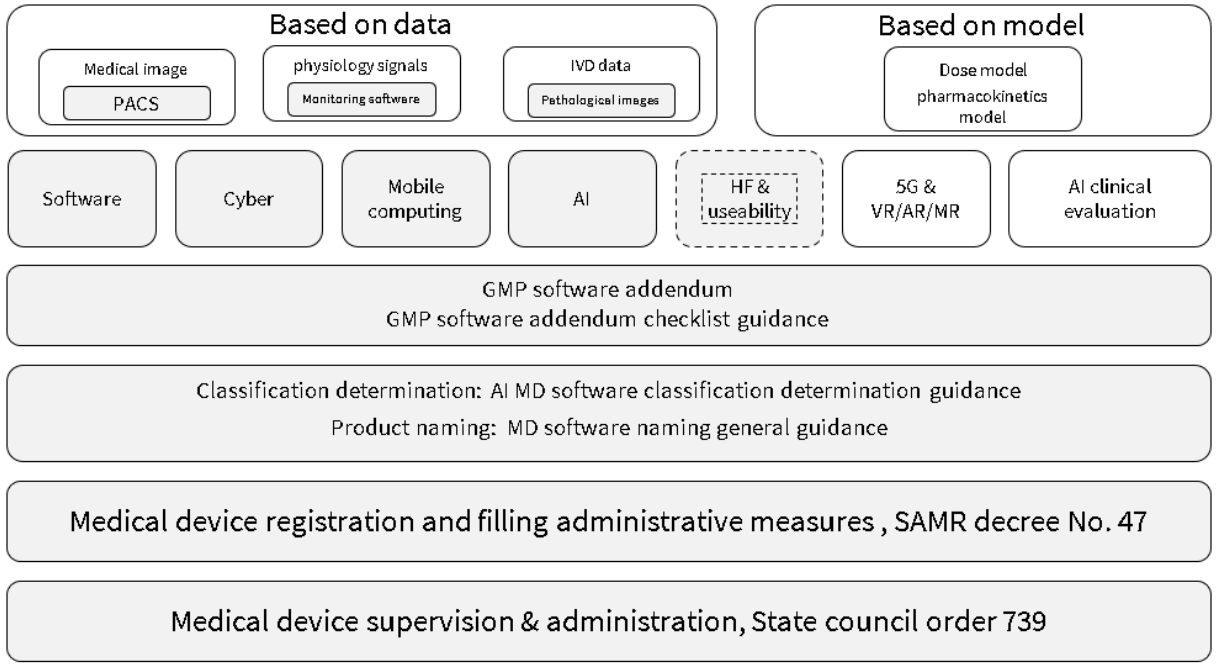

Setting the Standard | Harmonization

Continuing to press for standard based approaches

Harmonizing Around Standards

Continuing to push for the IMDRF approach

ML enabled MD definitions document published May 2022

No significant change in the past year

Currently under development

- N12 companion document to further develop the notion of risk in software product
- New work item on Good Machine Learning Principles (GMLP): objective to identify high level principles that different jurisdictions can agree on

Broader Acceptance of Pathways

Clarifying and expanding new pathways

1) Predictable reviews and requirements for locked AI/ML
2) Expansion (global) of PCCP clarity and acceptance
3) Clarity on unlocked/generative/adaptive AI/ML

AI SUMMIT
CINCINNATI, OH · NOVEMBER 14–16, 2023

Update on AI/ML Standards

Pat Baird, Philips

pat.baird@philips.com

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Inspiring Collaboration. Leading Innovation. Making a difference.

No Global Consensus on AI definition

- Warning! Different stakeholders in different regions have different definitions for AI/ML concepts.
- International Medical Device Regulators Forum (IMDRF) published a glossary in May 2022
  - IMDRF started with 16 different definitions
  - Eventually, IMDRF decide to not define AI at all
- Expect a lot of variation and feel free to get clarification from regulators, providers, standards organizations, and developers.
- Advice: when you disagree with someone, you might want to check to see if you both are talking about the same topic..

What is AI? FDA:

Artificial Intelligence has been broadly defined as the science and engineering of making intelligent machines, especially intelligent computer programs (McCarthy, 2007). Artificial intelligence can use different techniques, including models based on statistical analysis of data, expert systems that primarily rely on if-then statements, and machine learning.

Machine Learning is an artificial intelligence technique that can be used to design and train software algorithms to learn from and act on data. Software developers can use machine learning to create an algorithm that is 'locked' so that its function does not change, or 'adaptive' so its behavior can change over time based on new data.
Some real-world examples of artificial intelligence and machine learning technologies include:

- An imaging system that uses algorithms to give diagnostic information for skin cancer in patients.
- A smart sensor device that estimates the probability of a heart attack.

NOTE: AI encompasses wide spectrum of software tech ranging from simpler "if -> then" statements, look up tables to Machine Learning to Deep Learning.

What is AI? Physician Viewpoint..

"Artificial intelligence constitutes a host of computational methods that produce systems that perform tasks **normally requiring human intelligence**. These computational methods include, but are not limited to, machine image recognition, natural language processing, and machine learning. However, in health care a more appropriate term is "**augmented intelligence**" (AI), reflecting the enhanced capabilities of human clinical decision making when coupled with these computational methods and systems."
Augmented intelligence in health care (AMA Board of Trustees, 2018, emphasis **added**)

What is different about ML vs. traditional quality/regulatory?

When developing software to help with image recognition, diagnosis, identification of high-risk patients, etc., we usually talk to human experts to understand how they make their decisions, and we then write software to mimic their thought process.

For Machine *Learn*ing systems, we have a relatively generic algorithm that has the capability to learn, and we provide it with a lot of data, and it refines itself according to patterns that it sees in the data. **The data is what is defining the logic, not human experts.** But this leads to some problems..

- The data might be **incomplete** (e.g. missing important elements)
- The data might be **incorrect** (e.g. data from wearables might not be 100% accurate)
- Patterns may or may **not be relevant** (e.g. might be a coincidence in the data that is not clinically relevant);
- The data might not fully **represent** the target population
  - If you train using data from a hospital in Denver Colorado (Snow!) is it applicable to a retirement community in Palm Beach Florida (Sun!) ??
  - Certain demographics cannot afford healthcare, and therefore the software assumes that they are healthy since they don't visit the hospital very often.
- Health & Healthcare **change over time** (e.g. more cases of flu during the winter – what if your data is from spring & summer?)

Success Factor: Good Data Handling Practices

One challenge is that AI seems mysterious and magical, and people think we need a whole new way of thinking about it.

I propose that we handle data according to these rules:

- Keep records / retain information on the origin of the sample
- Sourcing, processing, preservation, testing and handling should be done in a safe manner
- Protect against contamination, viruses

Note: these concepts are already captured in IMDRF GRRP WGN47 FINAL:2018 document – when talking about tissue samples !!

My point is that we already know many good practices that simply need to be adapted for AI. We don't need to re-invent the wheel..

Image source: https://xkcd.com/1838/

Does the Data tell the story you think it is telling you?

I was talking to an owner of a McDonald's restaurant, and he had an interesting story about data. He was curious about how many customers ask for a cup of water. Normally, he can get a report from the cash registers about how much of a certain product is sold every month, but since water is free, it did not show up on the sales summary for the month. His restaurant didn't sell fish sandwiches, so he asked his employees to use the "Fish Sandwich" button whenever some asked for water.

After doing this for a few weeks, he started receiving a shipment of tartar-sauce packets from McDonald's supply – they automatically send shipments of things like napkins and paper cups and straws and sauce based on sales, and they saw that he was selling a lot of fish sandwiches so they sent him sauce to help him out.

How does this apply to ML? I know of several instances where the patient had a certain medical condition, but their insurance would not cover treatment. The physician lied on their medical record because the insurance would give treatment for this other condition.

My point is: If you are getting results that don't seem right, is it due to a security breach? Or is it a physician trying to help a patient in an unusual way? Or is there some other use of the data that you don't understand?

(BTW, you are going to remember this story the next time you go by McDonalds…)

Examples of Data without Knowledge

- There was a ML system trained to help pneumonia patients get the right kind of treatment, based on risk. Because patients with asthma are high risk when they get pneumonia, healthcare providers treated them more aggressively.
- This led the software to conclude that asthma patients must be low risk, because their survival rate is much higher.
- The software looked at data, without knowledge, and came to a conclusion that is opposite of reality.
- I was talking to a friend of mine, and she was looking at mortality rates for rare diseases at smaller hospital. As it turns out, the rate is much less at small hospitals than at larger hospitals! Surprised by this, she dug a little deeper, and realized that the smaller hospitals don't have the resources for these patients, so they are moved to larger hospitals.

Examples of Data without Knowledge

At a previous company, I was involved in product support and analysing postmarket data. One of the data sources was to look at sales of replacement parts to hospitals – sometimes as hospital would repair a device but not tell the manufacturer that something failed, so we looked at parts replacement orders as a quality indicator.

I noticed that a few hospitals had a sudden increase in buying parts in November and December. I thought it was perhaps related to the weather becoming colder in the winter – but the parts ordering was back to normal (or even lower) in January and February.

I reached out to one of these hospitals to ask what is causing the higher rate of failures – their response was that their yearly budget is “use it or lose it” – if there is any money left in their budget, they use that money to buy more parts, otherwise management would take the money back.

Challenge: Conflicting Needs

Data quality is a key point of pre-market reviews, is mentioned often in regulatory guidelines, and is the subject of many standards. One of the challenges that we see across the world is the conflicting needs of high quality data vs. ensuring data privacy. We might not be able to properly develop and tune ML models because of this conflict.

This could affect system performance by failing to identify factors that impact a patient's risk, because those factors have been removed from the data. Even if it is legal, some hospitals are reluctant to share their data with manufacturers due to security and privacy concerns.

There are also export restrictions on patient data which can make it difficult since ML software often runs on remote servers.

IMDRF

In September 2020, the IMDRF formed an AI working group, and the first project it to develop a glossary of AI/ML-related terms.

Member countries & industry were asked to submit draft definitions, and the WG debated whether or not a particular term is needed, as well as debating what the term actually means.

One very long discussion about scope – should the IMDRF efforts be looking at systems that use ML *anywhere* in the product, or only places where the ML impacts risk/benefit? E.g. ML is used to screen for breast cancer vs. ML is used to optimize workflow ?

Smaller group was formed to begin developing supporting text – let the larger WG debate definitions, let the smaller task force start writing an introduction, general discussion, warn readers that their definition of “validation” isn’t what the data scientists use, “supervised learning” doesn’t mean someone is looking over the shoulder of the AI to watch it learn, etc.

Current project is to get IMDRF alignment on GMLP


IMDRF International Medical Device Regulators Forum

Final Document

IMDRF/AIMD WG/N67

Machine Learning-enabled Medical Devices: Key Terms and Definitions

AUTHORING GROUP

Artificial Intelligence Medical Devices (AIMD) Working Group

Standards Overview

- ISO/IEC JTC1, SC42, developing horizontal standards for all industries. Many simultaneous projects and even more are being created. **Not likely that these horizontal standards would be required for medical devices, but they may contain ideas that we like and would carry to healthcare.**
- IEEE also developing a number of AI standards.
- Existing medical device committees are looking at this technology, and ISO/IEC TC215 has created a Task Force to help medical device people develop new standards – create a "landscape" of other standards, collect use cases, etc.
- CTA is developing general AI standards as well as healthcare-specific AI standards.
- AAMI & BSI have collaborated to develop healthcare AI standards.

My point is that there are many organizations looking at this technology and are committed to defining and sharing good practices.

Topics discussed in AI Standards

There will be a giant family of standards for ML systems, including:

- Definitions
- Governance
- Risk management
- Trustworthiness
- Security
- Managing Bias
- Verification & Validation
- Data Management
- Postmarket considerations

And others!

One of the major success factors will be keeping this at a manageable level.

Consumer Technology Association (CTA)

CTA is the trade association for the consumer technology industry (all consumer industries – not just healthcare)

AI standards committee (R13) & Health Care working group (R13 WG1) have published:

- "Definitions / Characteristics of AI in Health Care (ANSI/CTA-2089.1)"
- "The Use of AI in Health Care: Trustworthiness (ANSI/CTA-2090)"
- "The Use of Artificial Intelligence in Health Care: Managing, Characterizing and Safeguarding Data (ANSI/CTA-2107)"
- "Artificial Intelligence in Health Care: Practices for Identifying and Managing Bias (ANSI/CTA-2116)

The current project is a guide on a user-facing "Nutrition Label"

(BTW, ***CTA Get's Stuff Done!***)

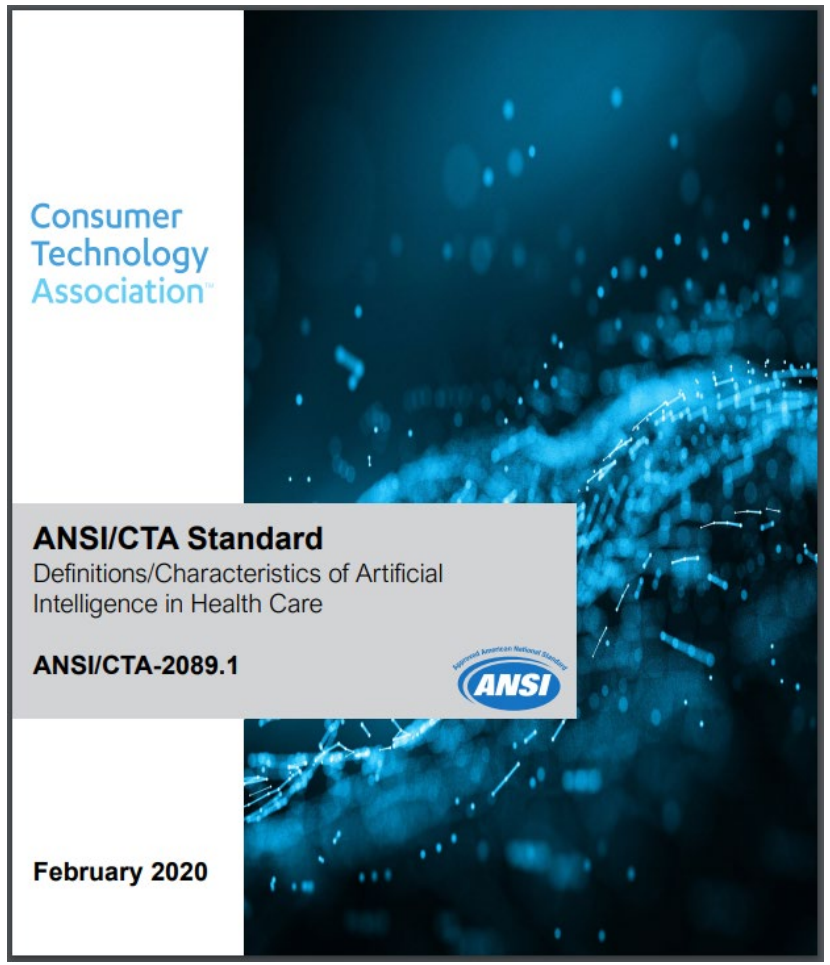

https://shop.cta.tech/collections/standards/artificial-intelligence

Current Project: IEC/TC62 PT 63450 AI-enabled Medical Devices – Methods for the Technical Verification and Validation

"This document establishes methods for medical device manufacturers to **verify and validate** artificial intelligence / machine learning-enabled medical devices (AI/ML-MD), i.e. medical devices that use artificial intelligence, in part or in whole, to achieve their intended medical purpose. This includes verification and validation activities for the **model** of the artificial intelligence as well as **selection**, **metrological characterization** and **management** of the **data sets**.

Such activities are implemented at **various stages of the medical device lifecycle**, especially including design control, monitoring and design change.

This document is also applicable to any hardware or software utilizing artificial intelligence that impacts the intended use of a medical device"

AAMI Artificial Intelligence

After publishing a few whitepapers with BSI, AAMI/BSI started working on another whitepaper regarding AI risk management (for medical devices, “risk = safety”.)

Feedback we received on the whitepaper was “why are you doing another whitepaper? A standard would be more useful…”

ISO 14971 is a commonly used risk management standard for medical devices. Our first AI/ML standards project was to use the process in 14971 and identify new ways that AI/ML will fail (e.g. the stories I mentioned earlier) and potential risk controls.

This is being taken to ISO/IEC; the kickoff meeting is in December, and we will add LLM-related risks.

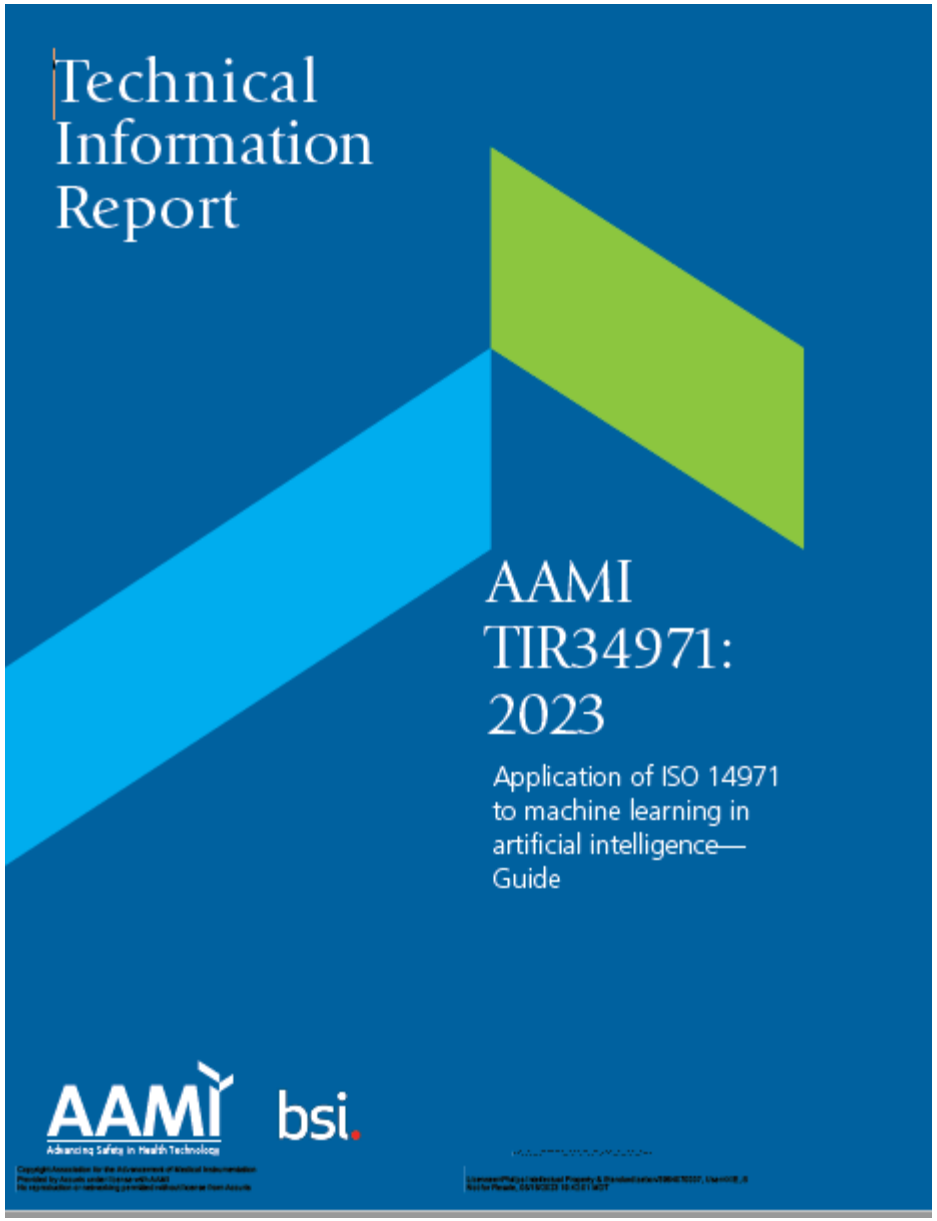

34971 ML-related hazards – Data Quality

- Incorrect data – the data has errors (eg corrupted data, data entry errors, calculation errors, inconsistent labelling of data.) This could also be related to how outliers are treated.
- Incomplete data – missing data (eg empty fields in a database)
- Subjective data – data that are influenced by beliefs or personal opinions rather than based on facts (e.g. patient reporting pain scales, differences in professional opinions)
- Inconsistent data – the quality of the data being used might vary from source to source, or even from the same source, taken at different times.
- Atypical data – the quality of the data during development might not represent the quality of the data in actual use (e.g. high resolution vs low resolution mammograms.)

34971 ML-related hazards – Data Storage/Security/Privacy

Data collected during the development, testing and improvement of an ML application could be a very attractive target for malicious actors.

- Privacy failures – information might be disclosed to unauthorized persons. Although the data can be anonymized, the anonymization process can fail. Additionally, personally identifiable information might contain critical information for the algorithm and anonymizing the data may destroy this critical information.
- Bias due to privacy – not all patients are willing to share their data and this can vary by patient demographics. For example, older patients might be reluctant to share their data, resulting in a bias towards younger populations.

More will come in the future ML-cybersecurity standard…

34971 ML-related hazards – Overtrust

- Over-confidence – the user trusts the system too much and believes it will work in all situations.
- Perceived risk – user might perceive the risks to be lower than they really are and are more likely to trust or delegate to ML
- User workload – people are busy and don't have time to stop and think about the application; a busy user is more likely to trust the software.
- Self-confidence – the user could defer to a product's "superior judgement"
- Variation in social trust – different user populations (including different professions, different cultures) have varying levels of trust and the developers might not be aware of these differences.
- User policies: company policies may put their trust in the ML software, forcing users to agree with the ML application.

Note that people may initially be skeptical of ML systems, but if it performs well, they will trust the system – perhaps even in situations where they shouldn't (e.g. California fire example.)

34971 ML-related hazards – Other Considerations

- Failure to act – the opposite of overtrust - the ML is ignored
- Data drift – things change over time (e.g. changes in patient demographics, disease prevalence, and medical practice standards of care.)
- Abundance of data but a lack of knowledge – This involves the need to truly understand what the data could mean. Example: software was created with the goal of identifying high-risk pneumonia patients, based on what other medical conditions the patients might have and the resulting mortality rates for similar patients. This software identified patients who had asthma as being low risk, which was surprising since medical practice considers these a high-risk population, and they commonly receive early interventions. As this discrepancy was investigated, it was discovered that because the interventions were successful, the mortality rate was low, and therefore the software did not flag asthma patients as being high risk user might perceive the risks to be lower than they really are and are more likely to trust or delegate to ML

34971 ML-related hazards – Potential Risk Controls

- Data quality controls (e.g. ensure data is complete, correct, consistent, represents the target population, agreement on any annotations/adjustments made by experts)
- Manage bias (a separate topic on its own..)
- Post-market monitoring for changes in performance
- User engagement in the design of the product (to find assumptions the development team might have had)
- Clear communication regarding need for user oversight and potential hand-off situations (e.g. self-driving cars might need your help, so don't climb into the backseat and play games on your phone. (Yes, this has happened))
- Model confirmation via independent methods

Takeaway: There is a strong need for good documentation. I know that it's not fun to spend your time, but we need it to support the product in the future

Cybersecurity & Standards

- Standards organizations are starting to recognize that this is an issue and have started some projects to address it.
- ISO/IEC SC27 & SC42 have joined forces to work on an international standard for security of ML systems; however, this is a horizontal standard across all sectors – I am hoping that we can have an informative annex that addresses some of the unique considerations that we have in healthcare.
- CTA has published a whitepaper and is currently working on standard – we will talk more about this is the cybersecurity session...

Free!
https://shop.cta.tech/products/cybersecurity-threats-and-security-controls-for-machine-learning-based-systems-pdf

Observations about AI

As you get involved in AI, there are some things that you notice that you would not have originally anticipated.

After a long day of getting nowhere in developing a standard for medical device interoperability, I commented “before we can get machines to talk to each other, we need to be better at getting people to talk to each other.”

This started me thinking about some of the ironies embedded in technology and in AI.

Lessons Learned from 1980s pop culture

I once asked a nurse "What has changed in the past 10 years?" and she replied, "I entered this field to take care of patients, but I spend all my time taking care of technology!"

This was 15 years ago. Things have not improved..

This was noted (by another industry) a long time ago -- in the 1980s,

the band Styx recorded the song "Mr. Roboto"; the lyrics include:

The problem is plain to see
Too much technology
Machines to save our lives
Machines de-humanize

Challenge: I am asking you "Renegades" with "Too Much Time on Your Hands" to stop "Fooling Yourself" and "Come Sail Away" with this "Blue Collar Man" on a mission to re-humanize healthcare. Are you with me??

AI SUMMIT
CINCINNATI, OH · NOVEMBER 14–16, 2023

Questions?

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Inspiring Collaboration. Leading Innovation. Making a difference.