Exponent®

Regulatory Changes for AI/ML Devices and Cybersecurity

Elizabeth Groves, PhD

Surya Sharma, PhD

Agenda

- Who Are We?
- Regulations Are Coming
- Foundation Models
- FDA Guidance on Software, AI/ML, and Cybersecurity

Who Are We?

Exponent is a multi-disciplinary engineering and scientific consulting firm that brings together more than 90 different disciplines to solve important **engineering, science, regulatory**, and **business** issues facing our clients.

Ex|50 Celebrating 50 Years

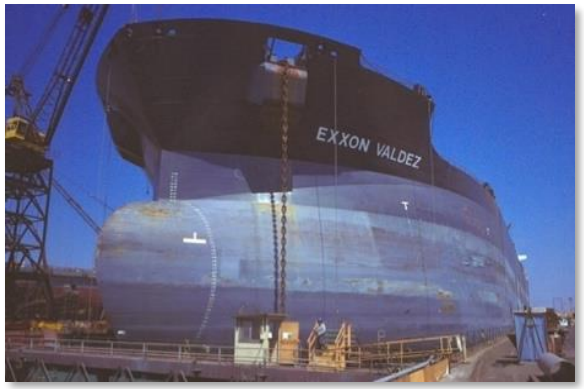

Our Commitment to the Advancement of Science

900+
Consulting Staff

500+
Doctoral-level professionals

35+
Exponent staff serve on editorial review boards or peer reviews

50+
Exponent staff teach or have advisory appointments at

35+
colleges and universities

75+
Exponent staff serve

65+
different organizations on

250+
Technical and engineering standards committees and scientific advisory boards

Global Engagement

Elizabeth Groves, Ph.D.
Managing Scientist
Data Sciences

Menlo Park
(650) 688-7147
egroves@exponent.com

Physicist | Software Developer | Data Scientist

Expert support in:
- Product Development & Failure
- FDA/regulatory environments
- IP support

Surya Sharma, Ph.D.
Managing Scientist
Electrical Engineering and Computer Science

Menlo Park
(650) 374-8686
ssharma@exponent.com

Computer Scientist | Embedded Systems

Expert support in:
- Machine Learning / Artificial Intelligence
- Failure Analysis
- Product Safety

Why Are We Here – ML in the Industry

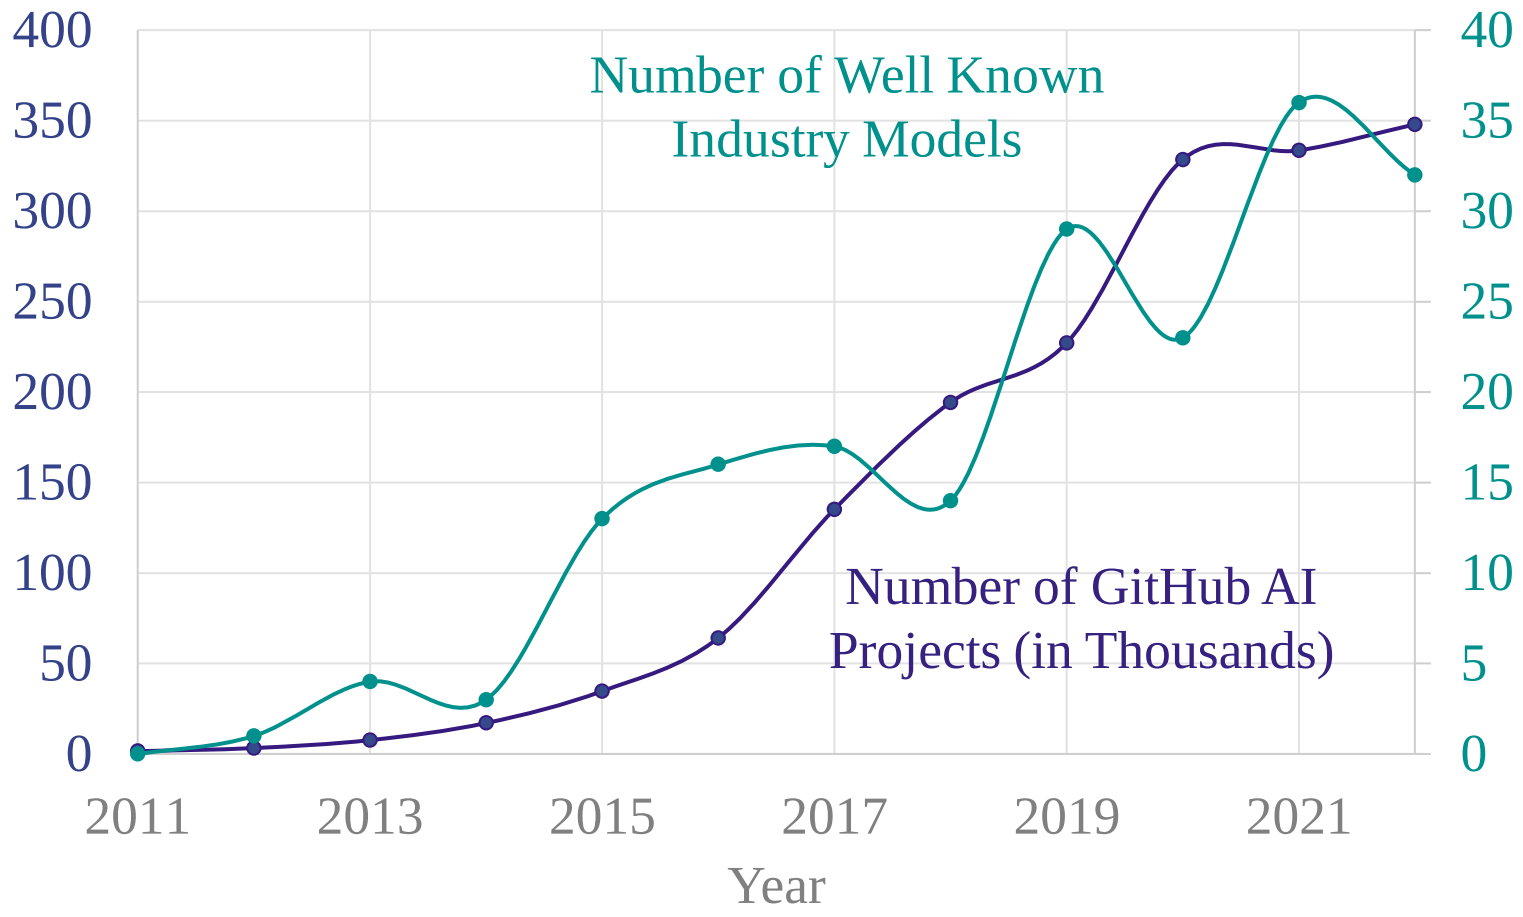

Source: GitHub, 2022; OECD.AI, 2022; Stanford 2023 AI Index Report

Why Are We Here – Medical Devices

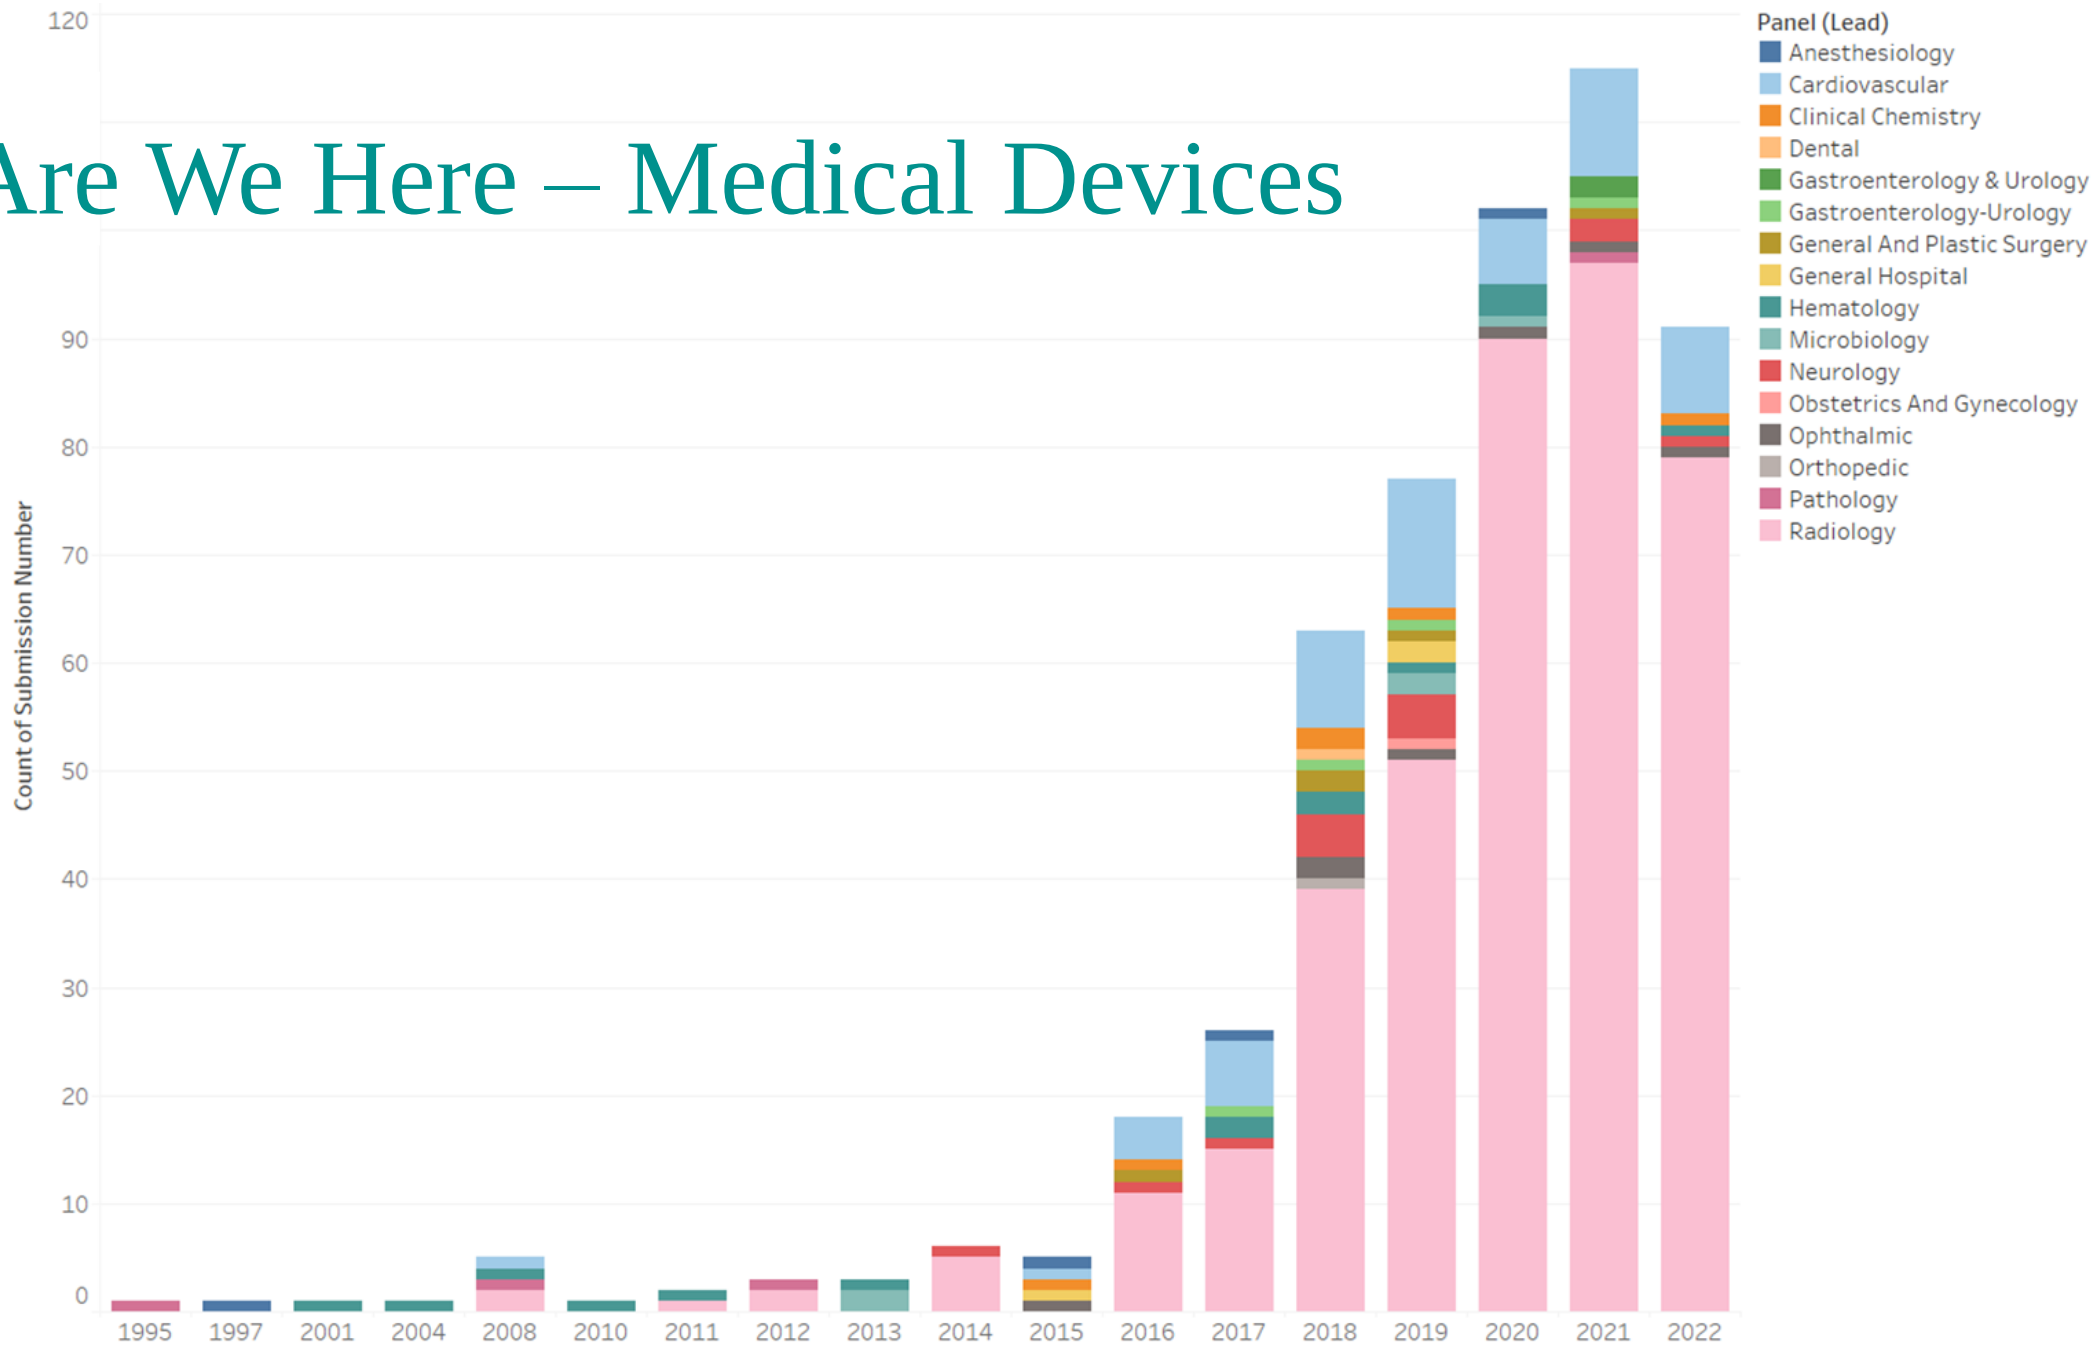

Why Are We Here - Regulations

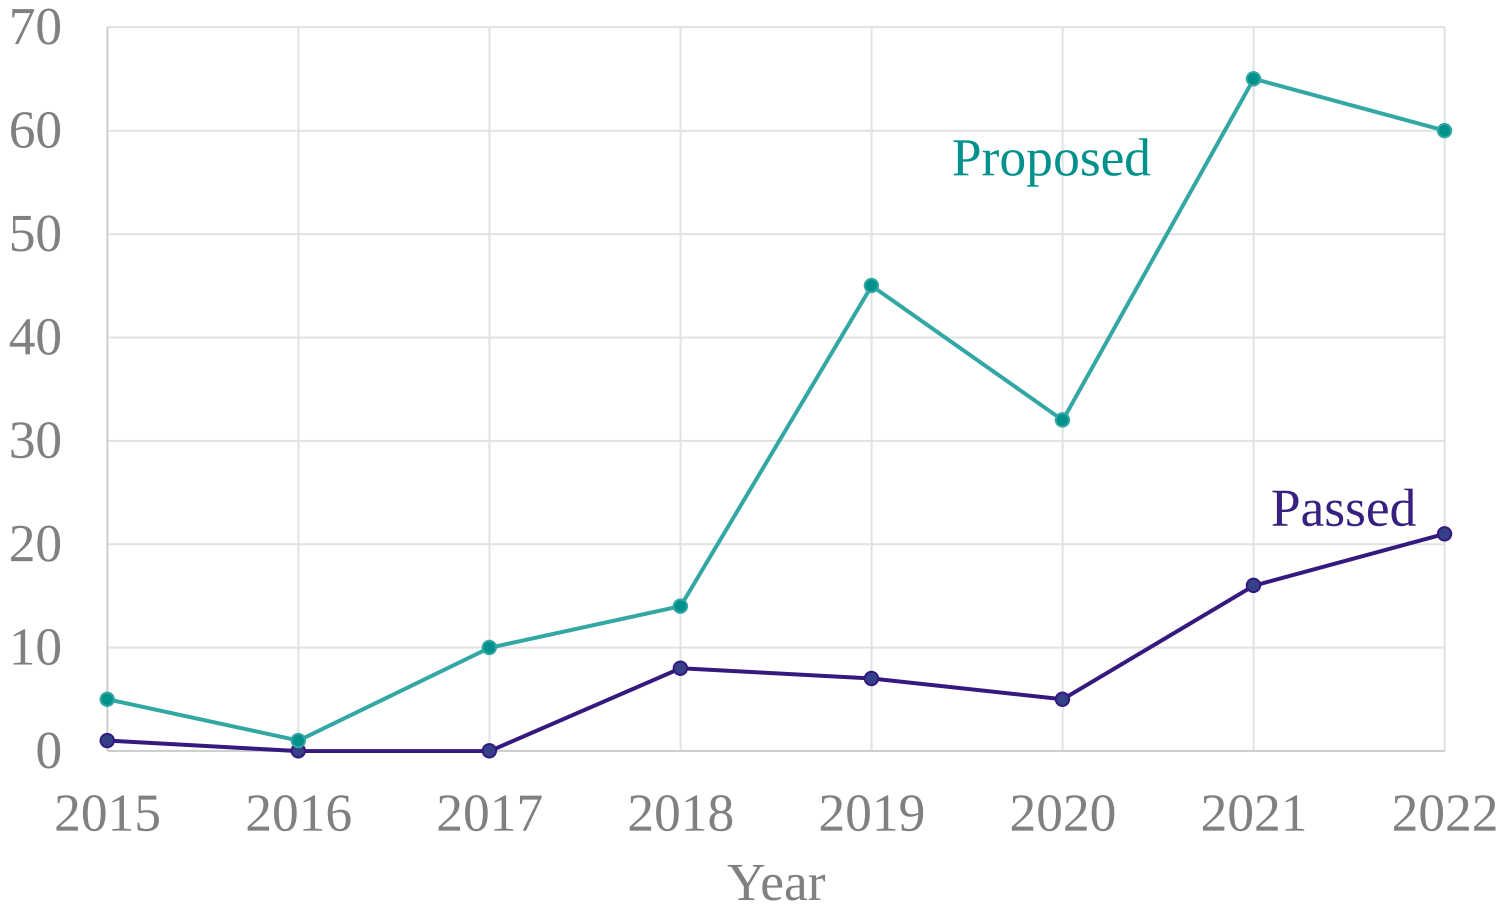

Source: Epoch, 2022; Stanford 2023 AI Index Report

The Washington Post

White House to unveil sweeping AI executive order next week

Tackling immigration and safety, the order would require advanced AI models to undergo assessments before they can be used by federal workers and ease barriers to entry for highly skilled workers

By Cat Zakrzewski, Cristiano Lima and Tyler Pager

Updated October 25, 2023 at 5:53 p.m. EDT | Published October 25, 2023 at 12:00 p.m. EDT

REUTERS®

Technology

Exclusive: G7 to agree AI code of conduct for companies

By **Foo Yun Chee**

October 29, 2023 6:41 AM PDT · Updated a month ago

FDA prioritizes guidance on AI, cybersecurity, pulse oximeters in stacked schedule for 2024

The administration is ramping up production of medtech guidance, adding 18 draft documents to the list of priorities for the upcoming financial year.

Published Oct. 19, 2023

Image generated by a text-to-image diffusion model.

What is Artificial Intelligence?

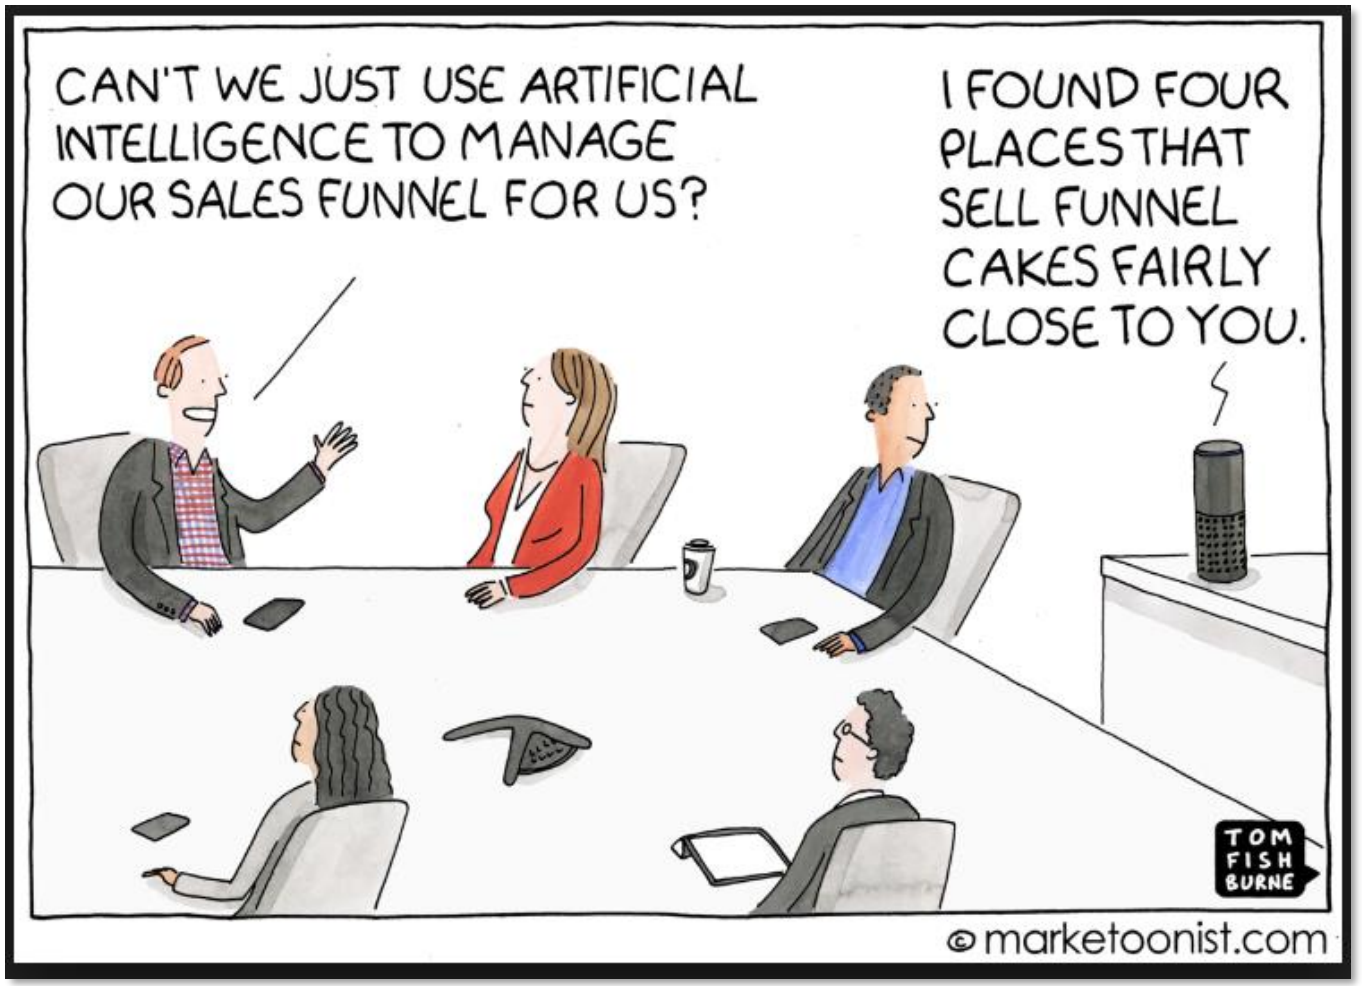

FDA Definition

Artificial Intelligence is the science and engineering of making intelligent machines, especially intelligent computer programs (McCarthy, 2007). AI 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.

- 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.

Interest in medical devices incorporating ML functionality has increased in recent years. Over the past decade, the FDA has reviewed and authorized a growing number of devices legally marketed (via 510(k) clearance, granted De Novo request, or approved PMA) with ML across many different fields of medicine—and expects this trend to continue.

October 19, 2023

Interest in medical devices incorporating AI/ML functionality has increased in recent years and even more so in recent months due to the development of large language models (LLMs). LLMs are AI models that are trained on very large datasets, enabling them to recognize, summarize, translate, predict, and generate content (for example: ChatGPT, Llama, Claude, PaLM, etc.). Over the past decade, the FDA has reviewed and authorized a growing number of devices (marketed via 510(k) clearance, granted De Novo request, or premarket approval) with AI/ML across many different fields of medicine—and expects this trend to continue. As of October 19, 2023, no device has been authorized that uses generative AI or artificial general intelligence (AGI) or is powered by large language models.

https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device#whatis

Recent EU and US Timeline

Since 2018, the European Union has published multiple communications and white papers on their approach to governing artificial intelligence

EU introduced their Artificial Intelligence Act

US releases an AI Bill of Rights

On Oct. 30, the Biden administration released a wide-ranging executive order.

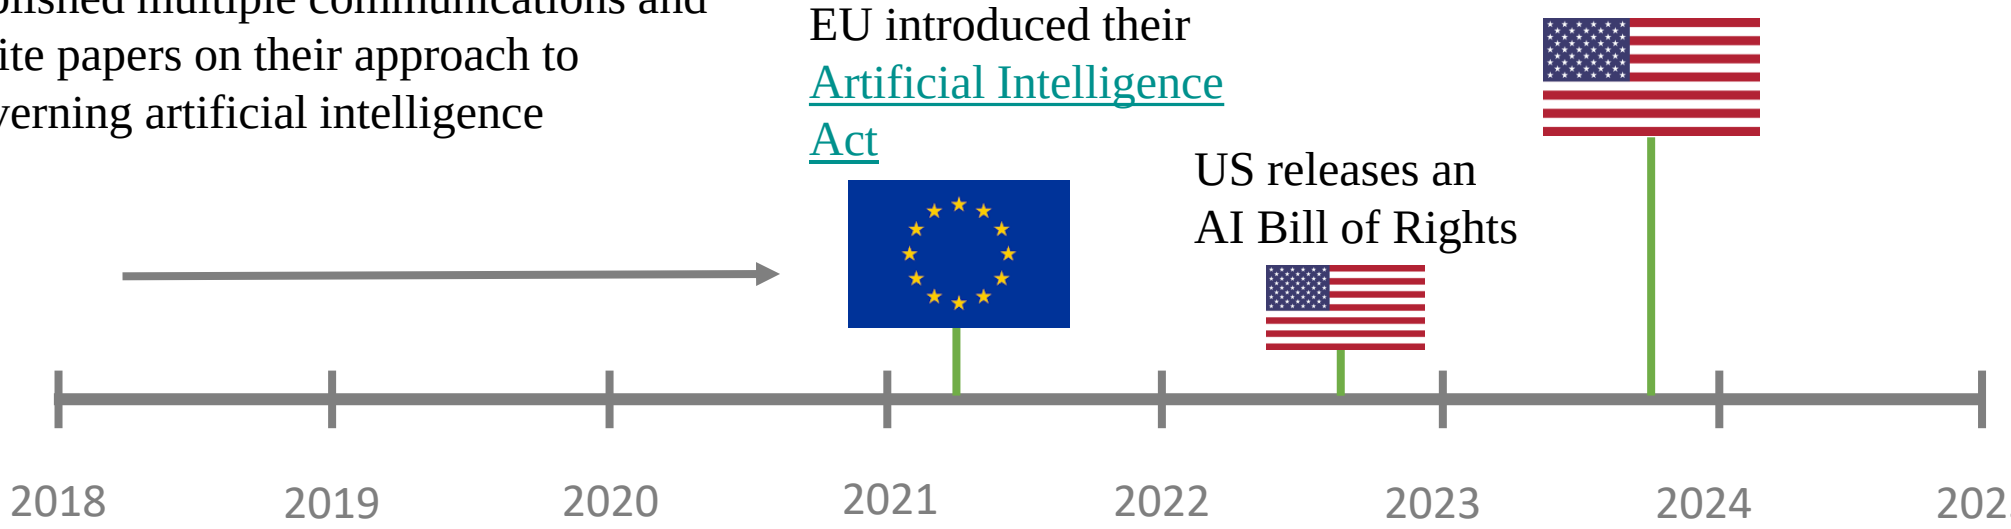

EU's AI Act can ban AI systems that are considered a threat carrying unacceptable risk.

The US President charges the Secretary of Defense, Secretary of Commerce, NIST, and others to establish norms, policies, and initiatives around AI, In some cases tasking specific deadlines such as 90 or 365 days from the date of the order.

Anticipating near- and long-term initiatives and changes to the regulatory landscape.

White House Executive Order on AI

- The new U.S. executive order is set to impact a wide range of domains, spanning from consumer products, housing and basic needs, transportation (including autonomous vehicles), through healthcare.
- Definition of AI signals impact for algorithms of varying levels of sophistication and architecture types, from supervised machine-learning predictive AI through generative AI.

"a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments."

- Discusses dual use foundation models.

	Guiding Principles
1	New standards for AI safety and security
2	Promoting innovation, competition, and collaboration
3	Commitment to supporting American workers
4	Advancing equity and civil rights while protecting against discrimination and abuse
5	Consumer protection laws and safeguards against fraud, bias, and other harms from AI
6	Protecting Americans' privacy and civil liberties
7	Ensuring responsible and effective government use of AI
8	Advancing American leadership abroad

What's a foundation model?

The mics were live and tape was rolling in the studio where the Miles Davis Quintet was recording dozens of tunes in 1956 for Prestige Records.

When an engineer asked for the next song's title, Davis shot back, "I'll play it, and tell you what it is later."

Like the prolific jazz trumpeter and composer, researchers have been generating AI models at a feverish pace, exploring new architectures and use cases. Focused on plowing new ground, they sometimes leave to others the job of categorizing their work.

A team of more than a hundred Stanford researchers collaborated to do just that in a 214-page paper released in the summer of 2021.

NVIDIA

Models that are trained on a broad set of unlabeled data that can be used for different tasks, with minimal fine-tuning (as opposed to task-specific models from the past decade).

IBM

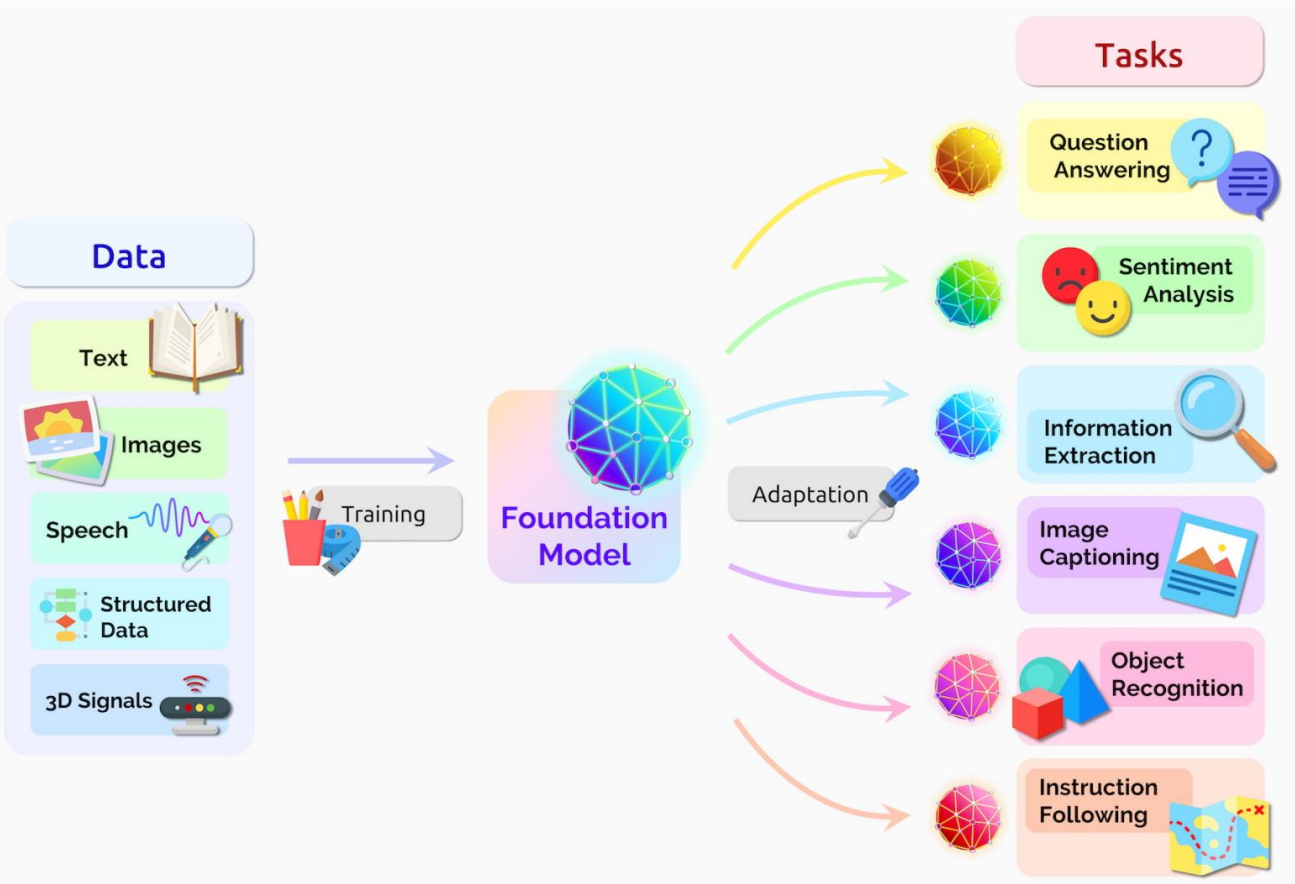

https://research.ibm.com/blog/what-are-foundation-models
https://blogs.nvidia.com/blog/what-are-foundation-models/
Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein et al. "On the opportunities and risks of foundation models." *arXiv preprint arXiv:2108.07258* (2021).

What's a dual use model?

- The general idea is that a technology, say, has an intended use or primary purpose which is good (or at least not bad) and a secondary purpose or use which is bad and is not intended by those who developed the technology in the first place.
  - Forge, J. A Note on the Definition of "Dual Use". *Sci Eng Ethics* **16**, 111–118 (2010). https://doi.org/10.1007/s11948-009-9159-9
- Research and technologies designed to generate benefits for civilians that can also be used for military purpose are termed "dual use".
  - Mahfoud, Tara, Christine Aicardi, Saheli Datta, and Nikolas Rose. "The limits of dual use." *Issues in Science and Technology* 34, no. 4 (2018): 73-78.
- Models that could be used for both beneficial or harmful machine learning based systems.
  - Henderson, Peter, Eric Mitchell, Christopher Manning, Dan Jurafsky, and Chelsea Finn. "Self-destructing models: Increasing the costs of harmful dual uses of foundation models." In Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, pp. 287-296. 2023.

Will this affect me?

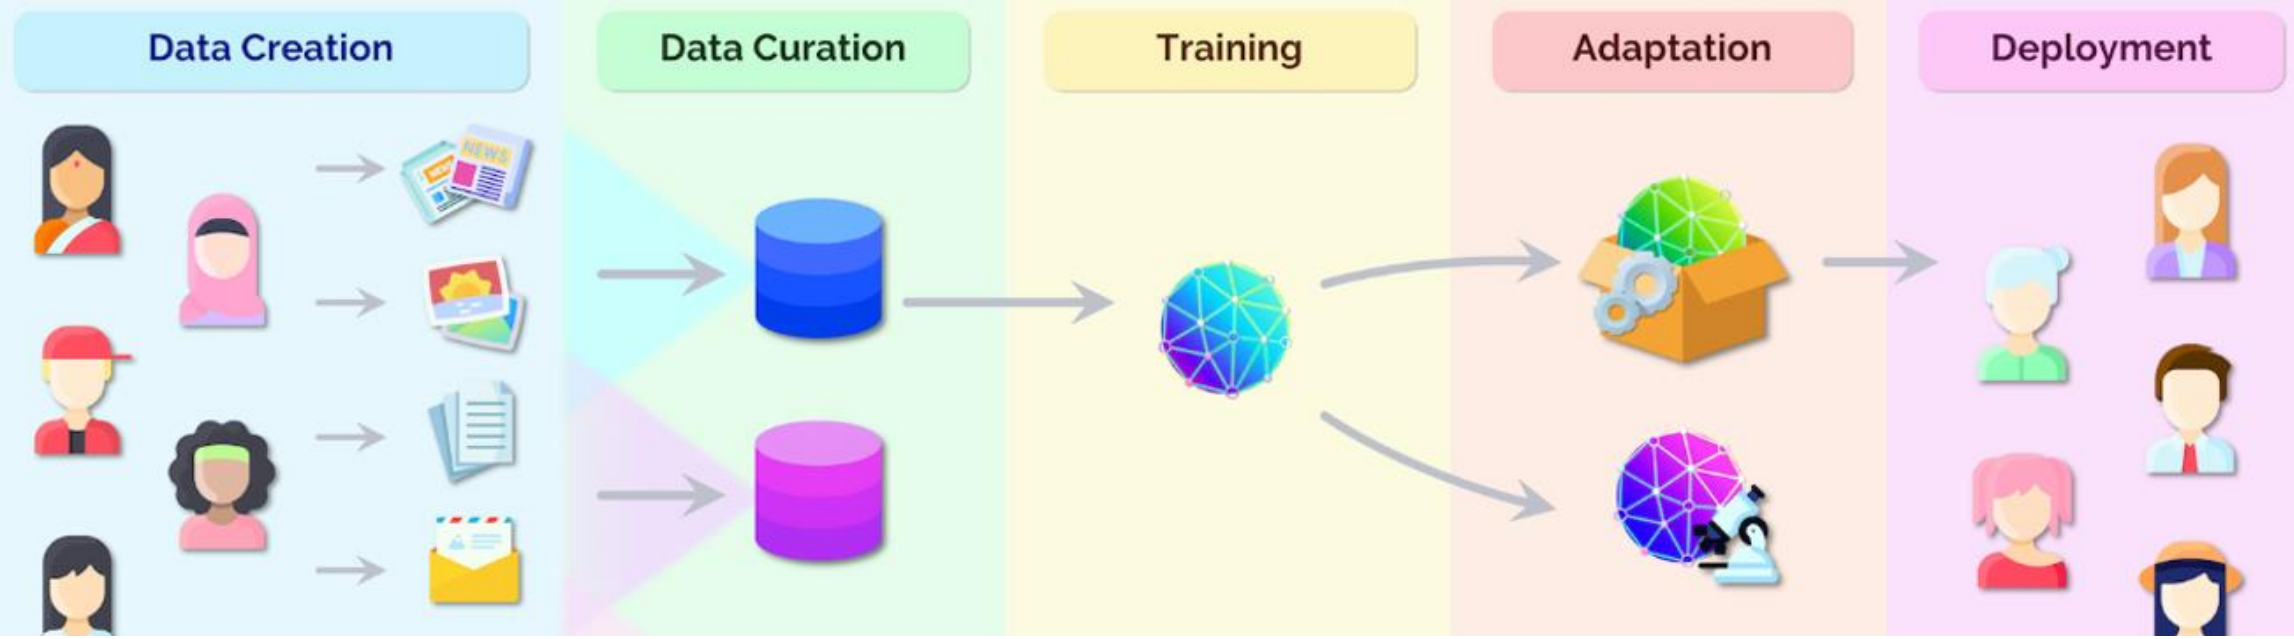

Do you build foundation or dual use models?

- Increased reporting burden.
- Increased burden on users of your models.
- Consider providing V&V and cybersecurity compliance data to your customers.

Do you use foundation or dual use models?

Is your product affected if

- There is an anomaly in the underlying model
- There is an update to the underlying model
- There is an attack on the underlying model

FDA Guidance relating to AI/ML and Cybersecurity

FDA Clearance Background

- Medical devices sold in the U.S. are subject to FDA clearance. One pathway is the 510(k) submission process.
- A 510(k) is a submission made pre-market to the FDA to demonstrate that the device to be marketed is at least as safe and effective, that is substantially equivalent, to an already marketed device in U.S.
- The FDA provides guidance on what documentation must be contained in a 510(k) submission and also publishes refusal to accept (RTA) guidance.
- Changes to software or device functionality can require a new 510(k). The FDA provides guidance: "Deciding When to Submit a 510(k) for a Software Change to an Existing Device"

Software Guidance (2023)

Software Documentation Elements	Basic Documentation Level	Enhanced Documentation Level
Documentation Level Evaluation	A statement indicating the Documentation Level and a description of the rationale for that level.	
Software Description	Software description, including overview of significant software features, functions, analyses, inputs, outputs, and hardware platforms.	
Risk Management File	Risk management plan, risk assessment demonstrating that risks have been appropriately mitigated, and risk management report.	
Software Requirements Specification (SRS)	SRS documentation, describing the needs or expectations for a system or software, presented in an organized format, at the software system level or subsystem level, as appropriate, and with sufficient information to understand the traceability of the information with respect to the other software documentation elements (e.g., risk management file, software design specification, system and software architecture design chart, software testing).	
System and Software Architecture Design	Detailed diagrams of the modules, layers, and interfaces that comprise the device, their relationships, the data inputs/outputs and flow of data, and how users or external products (including information technology (IT) infrastructure and peripherals) interact with the system and software.	

Software Guidance (2023) (cont.)

Software Documentation Elements	Basic Documentation Level	Enhanced Documentation Level
Software Design Specification (SDS)	FDA is not recommending the SDS as part of the premarket submission. Sponsor ***should*** document this information on the design via the DHF for the device. During premarket review, FDA may request additional information, if needed, to evaluate the safety and effectiveness of the device.	SDS documentation, including sufficient information that would allow FDA to understand the technical design details of how the software functions, how the software design completely and correctly implements all the requirements of the SRS, and how the software design traces to the SRS in terms of intended use, functionality, safety, and effectiveness.
Software Development, Configuration Management, and Maintenance Practices	A summary of the life cycle development plan and a summary of configuration management and maintenance activities; OR A Declaration of Conformity to the FDA-recognized version of IEC 62304.	Basic Documentation Level, PLUS complete configuration management and maintenance plan document(s); OR A Declaration of Conformity to the FDA-recognized version of IEC 62304.
Software Testing as Part of Verification and Validation	A summary description of the testing activities at the unit, integration and system levels; AND System level test protocol including expected results, observed results, pass/fail determination, and system level test report.	Basic Documentation Level, PLUS unit and integration level test protocols including expected results, observed results, pass/fail determination, and unit and integration level test reports.
Software Version History	A history of tested software versions including the date, version number, and a brief description of all changes relative to the previously tested software version.	
Unresolved Software Anomalies	List of remaining unresolved software anomalies with an evaluation of the impact of each unresolved software anomaly on the device's safety and effectiveness.	

Off-the-Shelf Software Guidance (2023)

Software Documentation Elements	Basic Documentation Level	Enhanced Documentation Level
Description of OTS Software	An overview and description of the OTS software and actions taken for the continued safe and effective use of the medical device.	
Risk Assessment of OTS Software	Risk assessment demonstrating that risks related to the use of OTS software have been appropriately mitigated.	
Software Testing as part of Verification and Validation	Test plans and results for the OTS software, commensurate with the Documentation Level (i.e., Basic or Enhanced) for the device.	
Assurance of Development Methodologies and Continued Maintenance of OTS Software	FDA is not recommending this documentation as part of the premarket submission. Sponsors ***should*** document this information via the design history file (DHF) for the device. During premarket review, FDA may request additional information, if needed, to evaluate the safety and effectiveness of the device.	Information to provide an assurance that the product development methodologies used by the OTS software developer are appropriate and sufficient, and mechanisms exist for assuring the continued performance, maintenance, and support of the OTS software

https://www.fda.gov/regulatory-information/search-fda-guidance-documents/shelf-software-use-medical-devices

Predetermined Change Control Plan (PCCP)

- Moving past "locked" medical algorithms, FDA draft guidance allows for some modifications through predetermined change control plans.
- Draft guidance addresses the adaptability of AI/ML-enabled medical devices.
- A PCCP consists of three components:
    1. Description of Modifications
    2. Modification Protocol
    3. Impact Assessment
- Sponsors should decide if they foresee future changes in their device, and if the administrative burden of a PCCP is worth the effort.

FDA Guidance relating to Cybersecurity

Date	Guidance	Summary
01/14/2005	Guidance for Industry – Cybersecurity for Networked Medical Devices Containing Off-the-Shelf Software	General principles that FDA considers to be applicable to software maintenance and cybersecurity for networked medical devices—specifically, those that incorporate off-the-shelf (OTS) software.
10/02/2014	Content of Premarket Submissions for Management of Cybersecurity in Medical Devices	Recommendations on information to include in FDA medical device premarket submissions for effective cybersecurity management*. Manufacturers should address cybersecurity during the design and development of the medical device, identify assets, threats, and vulnerabilities, determination risk levels and suitable mitigation strategies, assessment of residual risk and risk acceptance criteria.
12/28/2016	Postmarket Management of Cybersecurity in Medical Devices	Emphasizes that manufacturers should monitor, identify, and address cybersecurity vulnerabilities and exploits as part of their postmarket management of medical devices. Establishes a risk-based framework for assessing when changes to medical devices for cybersecurity vulnerabilities require reporting to the Agency. This assessment is based on an evaluation of the likelihood of exploit, the impact of exploitation on the device's safety and essential performance, and the severity of patient harm if exploited.
04/08/2022	Cybersecurity in Medical Devices	Draft - Quality System Considerations and Content of Premarket Submissions
12/29/2022	Food and Drug Omnibus Reform Act (FDORA) signed into law	Newly defined category "cyber device" that 1. includes software validated, installed or authorized by the sponsor 2. has the ability to connect to the internet 3. contains characteristics vulnerable to cybersecurity threats
03/30/2023	Refuse to Accept Policy for Cyber Devices and Related Systems Under Section 524B of the FD&C Act	The sponsor shall 1. submit a plan to monitor, identify, and address postmarket cybersecurity vulnerabilities and exploits, including coordinated vulnerability disclosure and related procedures; 2. design, develop, and maintain processes and procedures to provide a reasonable assurance that the device and related systems are cybersecure, and make available postmarket updates and patches to the device and related systems to address vulnerabilities 3. provide a software bill of materials, including commercial, open-source, and off-the-shelf software components; 4. demonstrate reasonable assurance that the device and related systems are cybersecure.
09/27/2023	Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions	Detailed recommendations on conducting cybersecurity risk assessments, interoperability considerations, and documents to be included in premarket submissions to FDA. FDA recommends implementation and adoption of a "Secure Product Development Framework" or "SPDF", which consists of security risk management, security architecture, and cybersecurity testing. The risk management section also recommends Interoperability Considerations

* Effective cybersecurity management is intended to reduce the risk to patients by decreasing the likelihood that device functionality is intentionally or unintentionally compromised by inadequate cybersecurity.

What is a Cyber Device?

Image generated by a text-to-image diffusion model.

FDA Definition – “Cyber Device”

Section 524B(c) of the FD&C Act defines "cyber device" as a device that:

(1) includes ***software validated, installed, or authorized*** by the sponsor as a device or in a device,

(2) has the ***ability to connect to the internet***, and

(3) contains ***any*** such ***technological characteristics*** validated, installed, or authorized by the sponsor that could be ***vulnerable to the cybersecurity threats***.

If manufacturers are unsure as to whether their device is a cyber device, they may contact the FDA.

https://www.fda.gov/medical-devices/digital-health-center-excellence/cybersecurity-medical-devices-frequently-asked-questions-faqs

Cybersecurity in Medical Devices: Quality System Considerations (2023)

- Adoption of a “Secure Product Development Framework” or “SPDF.”
  - Threat Modeling
  - Cybersecurity Risk Assessment
  - Interoperability Considerations
  - Third-Party Software Components
  - Software Bill of Materials (SBOM)
  - Security Assessment of Unresolved Anomalies
  - TPLC Security Risk Management
- Designing for Security
- Transparency
- RTA (Refuse To Accept) Policy

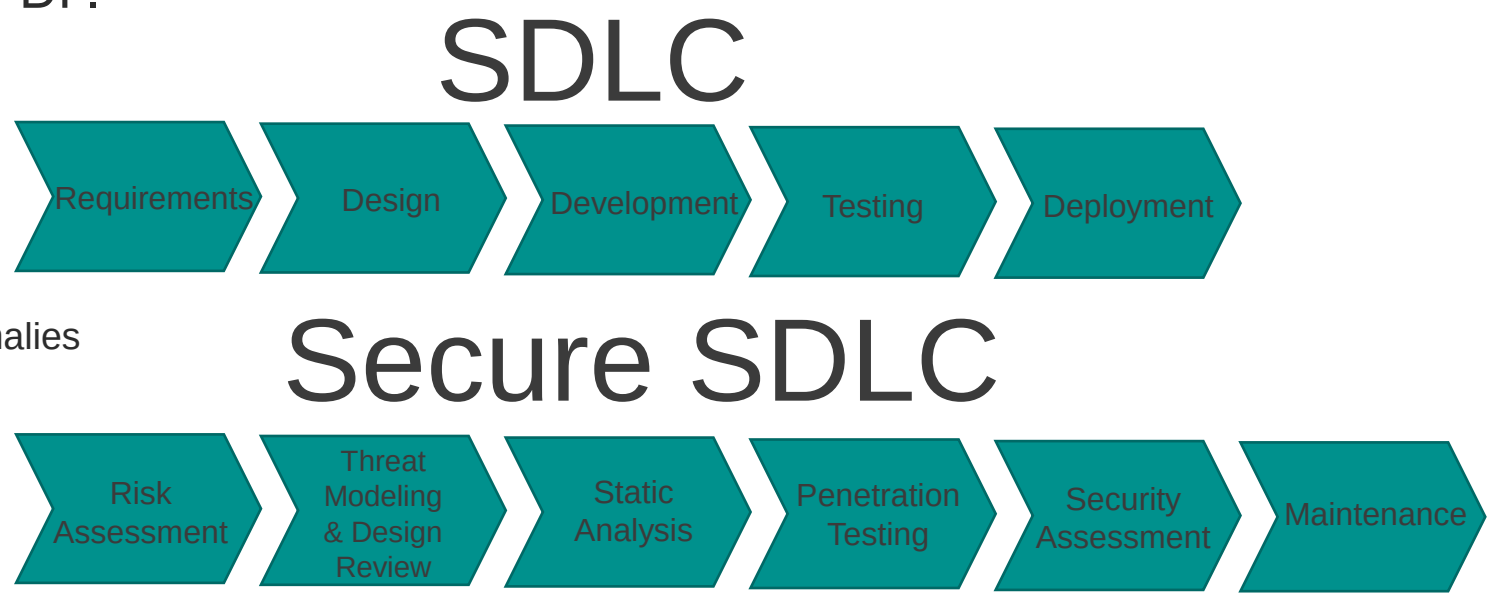

Example FDA Warning Letter

WARNING LETTER

May [redacted] **CMS#** [redacted]

[redacted] Inc.
[redacted]
[redacted], CA 9[redacted]

Dear [redacted]

During an inspection of your firm located at [redacted] on July 25 through August 12, 2022, investigators from the United States Food and Drug Administration (FDA) determined that your firm is a medical device manufacturer of the [redacted] System. Under section 201(h) of the Federal Food, Drug, and Cosmetic Act (the Act), 21 U.S.C. § 321(h), this product is a device because it is intended for use in the diagnosis of disease or other conditions or in the cure, mitigation, treatment, or prevention of disease, or to affect the structure or any function of the body. As stated in the [redacted] System indications for use, the system is intended to continuously record and report patient symptomatic and asymptomatic cardiac events and continuous electrocardiogram (ECG) information.

⋮

Unapproved Device Violations

You currently have clearance for the [redacted] under [redacted] for the following indications:

The [redacted] System is intended to capture, analyze and report symptomatic and asymptomatic cardiac events and continuous electrocardiogram (ECG) information for long-term monitoring. While continuously recording patient ECG, both patient-triggered and automatically detected arrhythmia events are transmitted to a monitoring center for reporting. After wear, a final report is generated based on beat-to-

⋮

[redacted] The reports are provided for review by the intended user to render a diagnosis based on clinical judgment and experience. It is not intended for use on critical care patients.

Thus, your device was cleared under [redacted] for long-term monitoring of arrhythmia events for non-critical care patients where real-time monitoring is not needed as reporting timeliness is not consistent with life-threatening arrhythmias. However, your marketing materials and other documentation, such as the document titled "[redacted]" and your website [redacted] monitors; last accessed February 28, 2023), state that the [redacted] System is intended for "near real-time monitoring" as a "mobile cardiac telemetry monitor," can provide notifications "immediately," and that it is intended for "high-risk patients." The claim that the device is intended as a mobile cardiac telemetry monitor implies this device is intended for high-risk patients and near real-time monitoring.[1] [redacted]

Our inspection also revealed that your firm made changes to the device without submission of a new 510(k). In response to our inspection, your firm provided more detail on the changes made to the device. Based on the information in your responses, the following changes were made to the device that require a new 510(k) submission:

a. Your document titled "[redacted] Design Changes, Brief Summary" and "**(b)(4)** Hardware and Firmware Design Changes" identifies hardware and firmware changes that raise new risks including, but not limited to, changes in the **(b)(4)**, and updating **(b)(4)** of your device including the **(b)(4)**. Specifically, these changes can affect the **(b)(4)** of the device (i.e., device **(b)(4)** and device **(b)(4)**) that would require new testing to support your assertion that the device's safety and effectiveness is unaffected. These changes can result in potential **(b)(4)** and may affect the basic safety of the device. According to Table B (technology, engineering, and performance changes) in the 510(k) Modifications Guidance, Question B5.2 should be answered "yes" because these changes raised new or significantly modified risks. Because these changes would require new **(b)(4)** testing and could significantly affect safety or effectiveness, a new 510(k) must be submitted for this change.

b. "[redacted] Design Changes, Brief Summary" describes a change to your device made in September 2019, "Improve **(b)(4)**: the algorithm in the **(b)(4)** was adjusted based on available data to improve **(b)(4)** performance for **(b)(4)** events that meet **(b)(4)** criteria. There were no changes to the algorithm that would significantly affect any clinical outcome." Changes to the **(b)(4)** algorithm can impact the accuracy of the detected arrhythmias, and testing would be needed to support that the algorithm's performance is unaffected by the change. A change in the **(b)(4)** for the algorithm could significantly impact safety or effectiveness, such as through missed or incorrect detection of events or arrhythmias, which could lead to patient injury due to lack of treatment. Question B5 in the 510(k) Modifications Guidance should have been answered "yes" because the change was to performance specifications. Subsequently, question B5.3 should have been answered "yes," because the modification to the algorithm necessitated clinical validation data in order to ensure the performance of the device was maintained, using relevant device-specific data from patients. Therefore, a new 510(k) must be submitted.

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/compliance-actions-and-activities/warning-letters

Thank You
Questions?