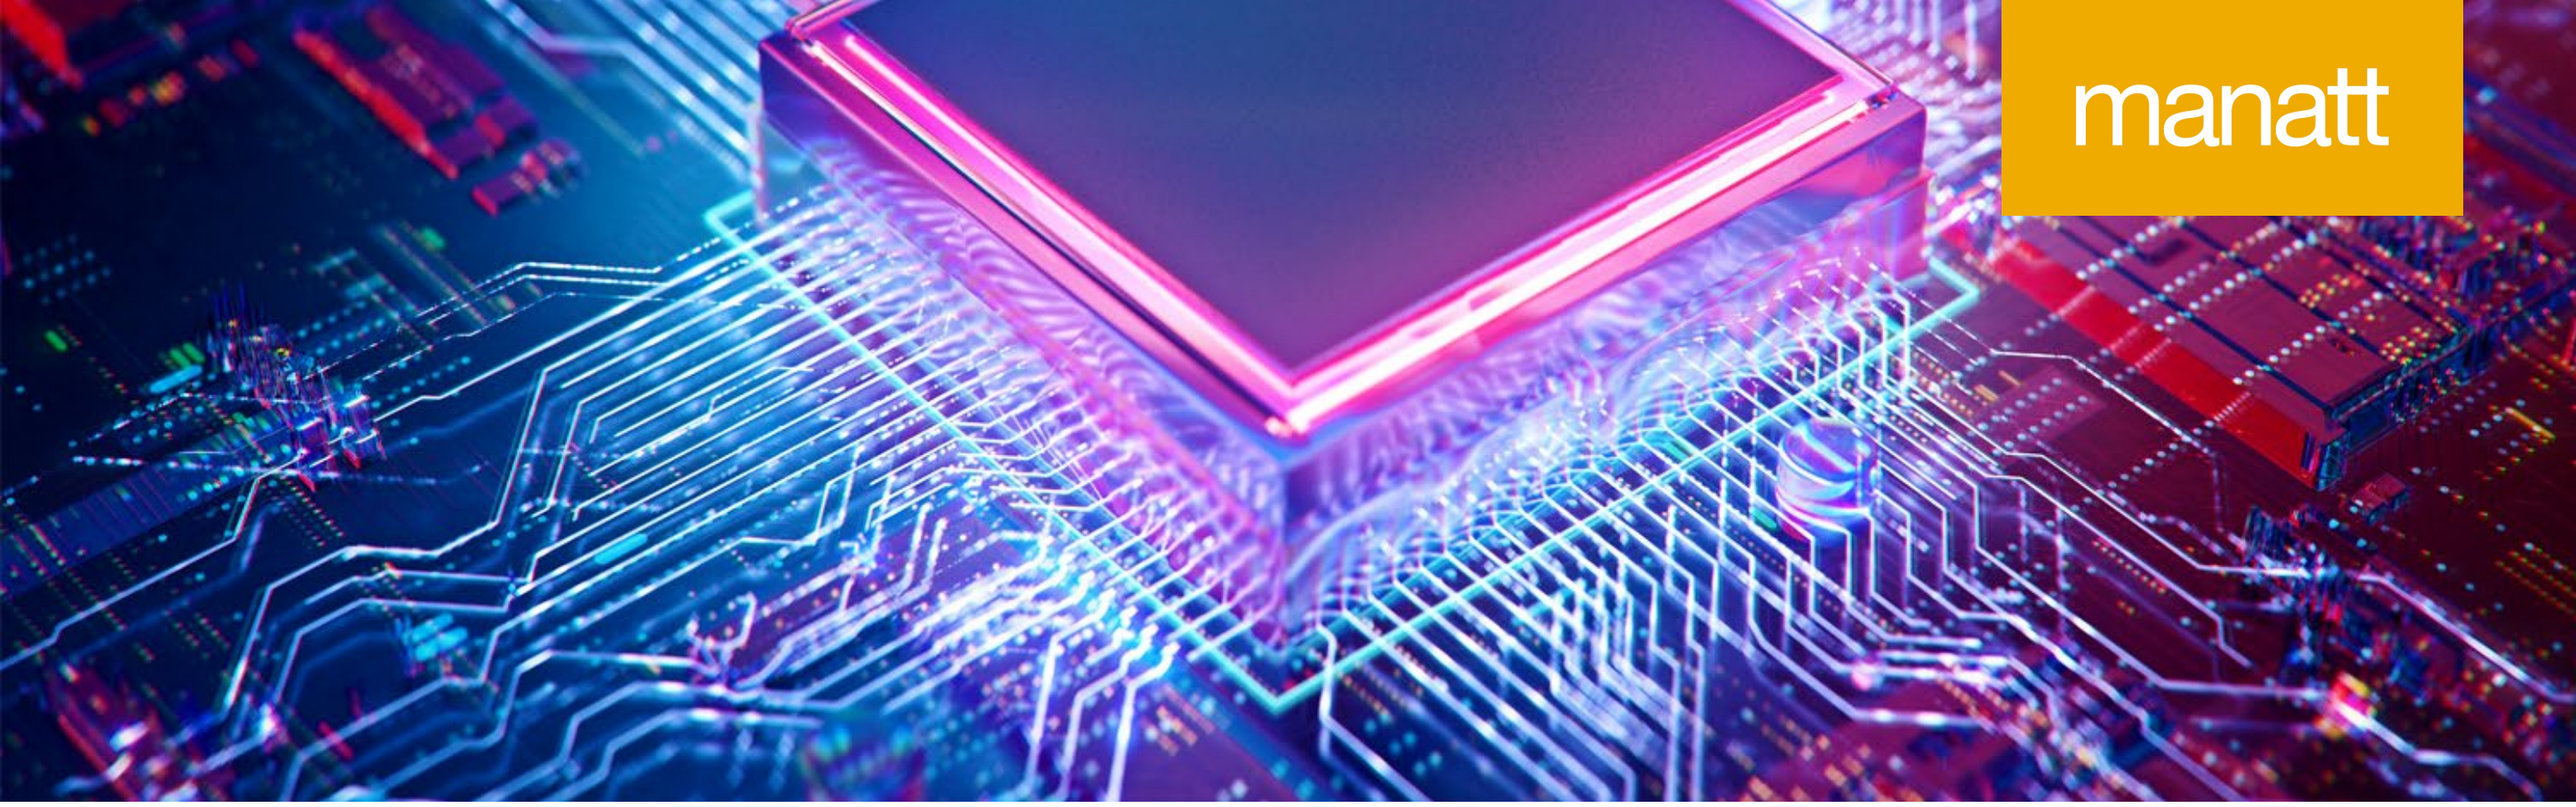

manatt

Health AI Primer – for CHCANYS and CHCs

March, 2024

Health AI – Progress and Opportunity

manatt

Some Definitions

AI enables computers to solve problems and perform tasks that would otherwise require human intelligence

- AI Has been around for 70+ years
- Humans control the operation and application of AI
- Data drives AI
- AI predicts future behavior based on statistics/past data, but does not conclude
- Generally, AI cannot think outside the data it is fed and has learned from

To be considered AI, the system must exhibit a level of learning and adapting.

- *AI is Not Automation.* Although both AI and automation rely on data and share a goal of streamlining tasks, the two concepts are not synonymous.

History of Artificial Intelligence (AI)

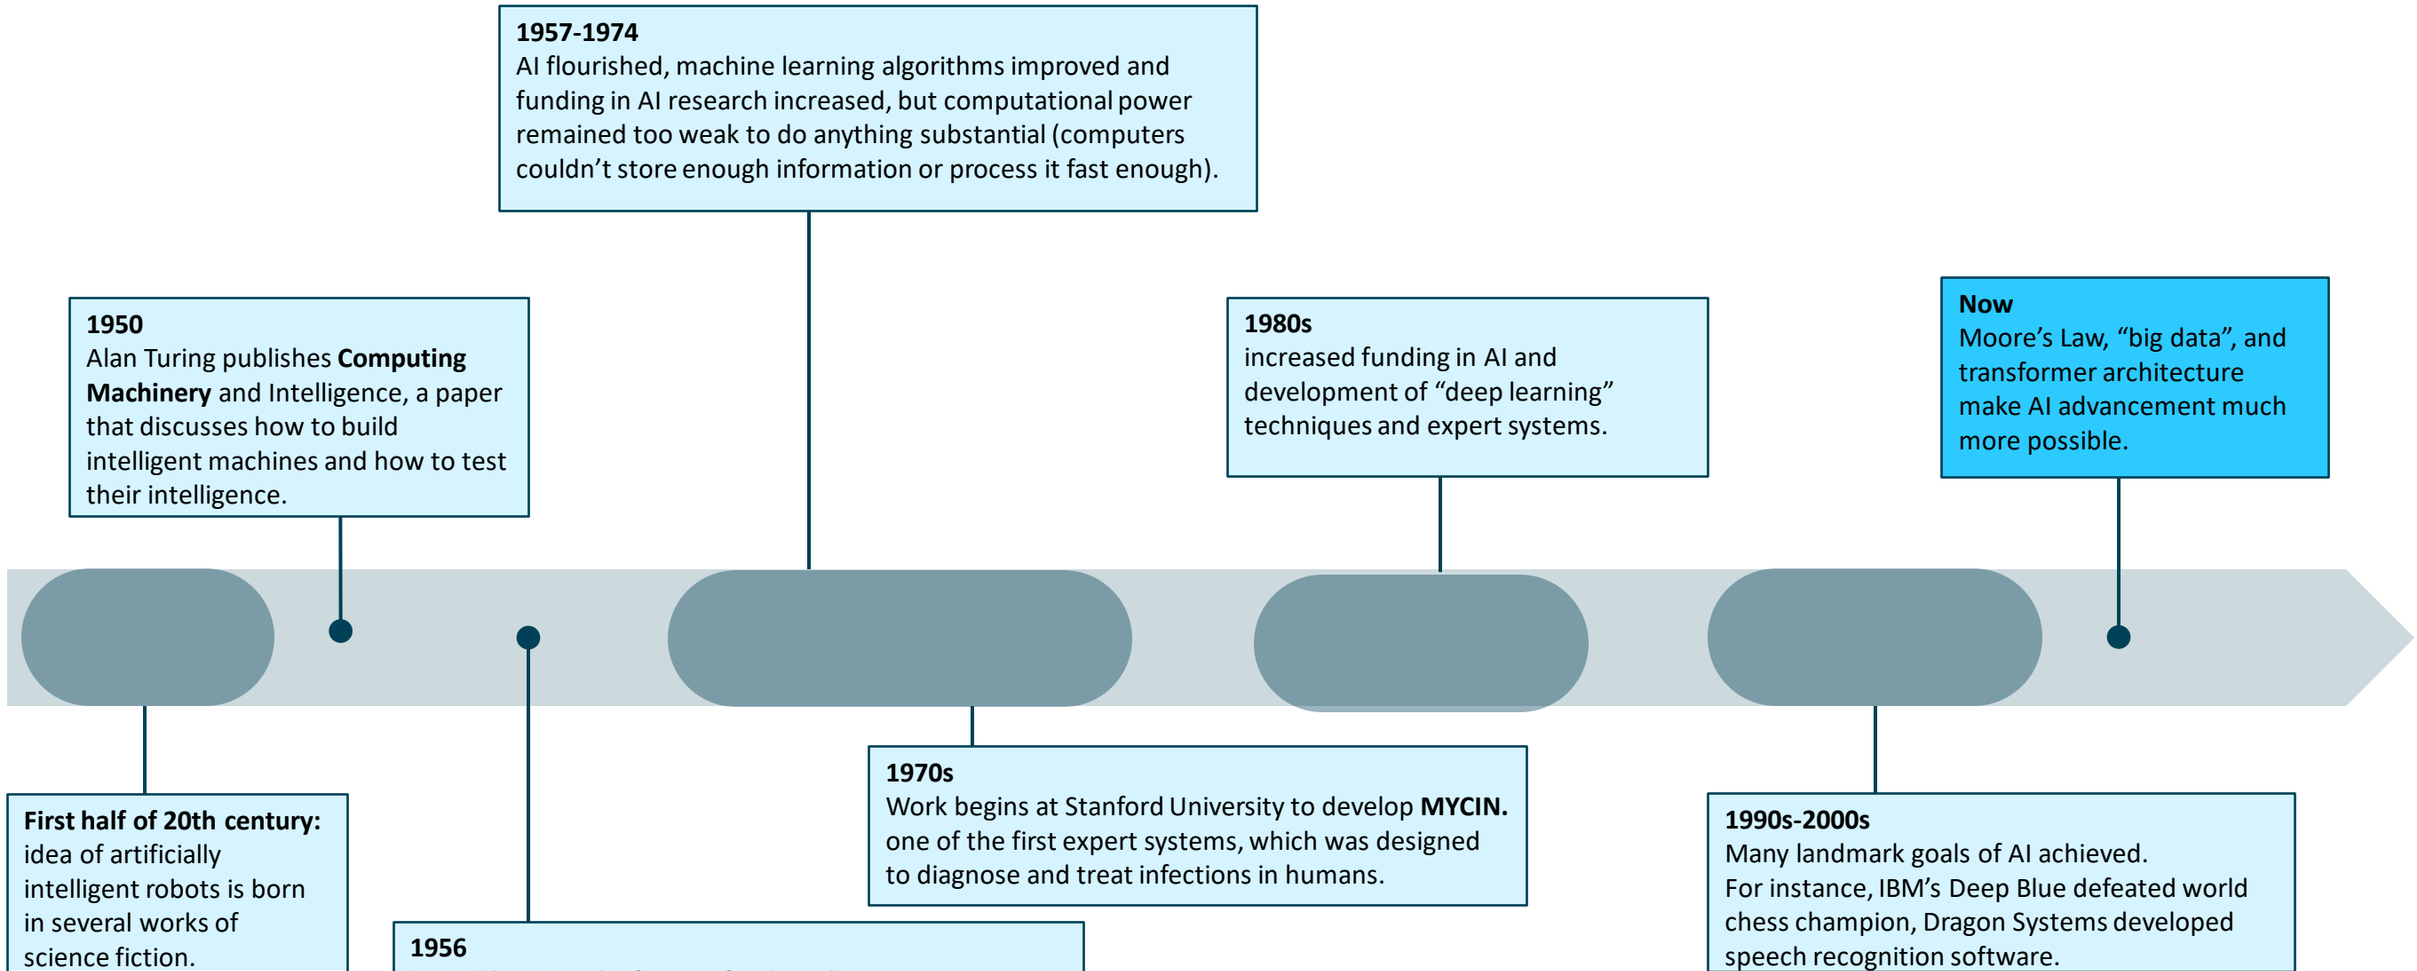

manatt

Types of AI

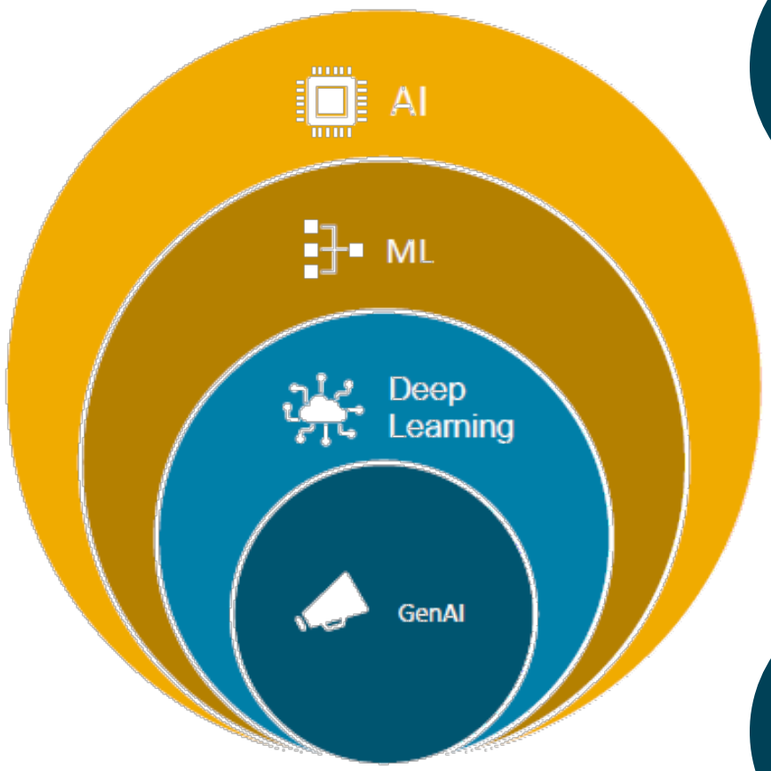

Machine Learning

A type of AI allows software-driven machines to learn on their own and improve performance with that information (computers identify, aggregate and derive meaning from data).

Deep Learning

A type of machine learning that involves passing data through multiple computer nodes that mimic the neural network of the human brain and enables the machine to interpret future data more accurately.

Generative AI

A type of deep learning model that takes raw data and "learns" to generate statistically probable outputs when prompted. Generative AI trains a computer model to recognize patterns within large language models.

AI Definitions; What is generative AI?; An Introduction to AI for Federal Judges

Machine Learning Training Techniques

Figure 3: Machine learning training techniques

Supervised learning

- Learns known patterns
- Takes labeled input data
- Predicts outcome/future

Unsupervised learning

- Learns unknown patterns
- Takes unlabeled input data
- Finds hidden patterns

Reinforcement learning

- Generates data
- Takes labeled input data
- Interacts with environment
- Learns series of actions

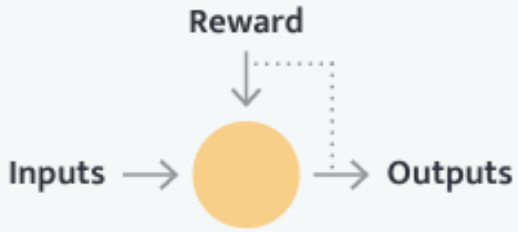

Source: American Medical Association. Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care. 2024.

Some Other Definitions for GenerativeAI (GenAI) and Augmented Intelligence (AuI)

GenAI allows computers to generate new content, revolutionizing the application of AI. Advances in GenAI and the rise in its popularity is underscored by the launch of ChatGPT (generative pre-trained transformer) in 2022.

Foundational Model

A type of AI-based model in which a model is trained on a large amount of (typically unlabeled) data which can be used for many applications. E.g., ChatGPT, DALL-E, Bard.

Large Language Model

A type of foundational model which learns the probabilities of occurrence of sequences of words from a corpus of text, whose probabilities are learned using textual corpora with trillions of words such that the resulting model has billions of parameters.

Augmented Intelligence (AuI)

A subset of artificial intelligence (AI) that uses machine learning to help humans make better decisions. The goal of augmented intelligence is to enhance human intelligence and decision-making, rather than replace it.

What are foundation models?; Creation and Adoption of LLM in Medicine; What Is Augmented Intelligence?

What is Different Now – AI Models in the Past and Present

Figure. Artificial Intelligence (AI) 1.0, 2.0, and 3.0

Approximate beginning year	1950s **AI 1.0** Symbolic AI and probabilistic models	2011 **AI 2.0** Deep learning	2018-2022 **AI 3.0** Foundation models
Core functionality and key features	Follows directly encoded rules (if-then rules or decision trees)	Predicts and/or classifies information Task-specific (1 task at a time); requires new data and retraining to perform new tasks	Generates new content (text, sound, images) Performs different types of tasks without new data or retraining; prompt creates new model behaviors
Training method	Rules based on expert knowledge are hand-encoded in traditional programming	Learning patterns based on examples labeled as ground truth	Self-supervised learning from large datasets to predict the next word or sentence in a sequence
Performance capabilities	Follows decision path encoded in its rules. *Eg, ask a series of questions to determine whether a picture is a cat or a dog.*	Classifies information based on training: *"Is this a cat or a dog?"* *"How many dogs will be in the park at noon?"*	Interprets and responds to complex questions: *"Explain the difference between a cat and a dog."*
Examples of performance	IBM's Deep Blue beat the world champion in chess **Health care:** Rule-based clinical decision support tools	Photo searching without manual tagging, voice recognition, language translation **Health care:** diabetic retinopathy detection, breast cancer and lung cancer screening, skin condition classification, predictions based on electronic health records	Writing assistants in word processors, software coding assistants, chatbots **Health care:** Med-PaLM and Med-PaLM-2, medically tuned large language models, PubMedGPT, BioGPT
Examples of challenges and risks	Human logic errors and bias in encoded rules lead to limited capability with real-world situations	Out-of-distribution problems (real-time data differs from training data) Catastrophic forgetting (not remembering early parts of a long sequence of text) Bias related to underlying training data	Hallucinations (plausible but incorrect responses based solely on predictions) Grounding and attribution Bias related to underlying training data and semantics of language in datasets

Source: M Howell, G Corrado, K DeSalvo, "Three Epochs of Artificial Intelligence in Health Care," JAMA (Jan 2024)

manatt

Generative AI – Simple Version

Data → *Training* → **Foundational Model** → **Use the Model**

CommonCrawl (billions of webpages), books, academic articles, movie scripts, customer support dialogues, etc.

PaLM 2

BARD

ChatGPT
Google Bard
MedPaLM 2
Gemini
LLaMA
Etc.

- Prompting and Prompt Engineering
- Fine Tune to Improve Performance for a Specific Task
- Embed into Larger Data Architecture

manatt

Rapid Performance Progress: USMLE Accuracy

Performance of LLMs is advancing rapidly. The latest models have 85+% accuracy on the US Medical Licensing Exam, up from 33% only four years ago.

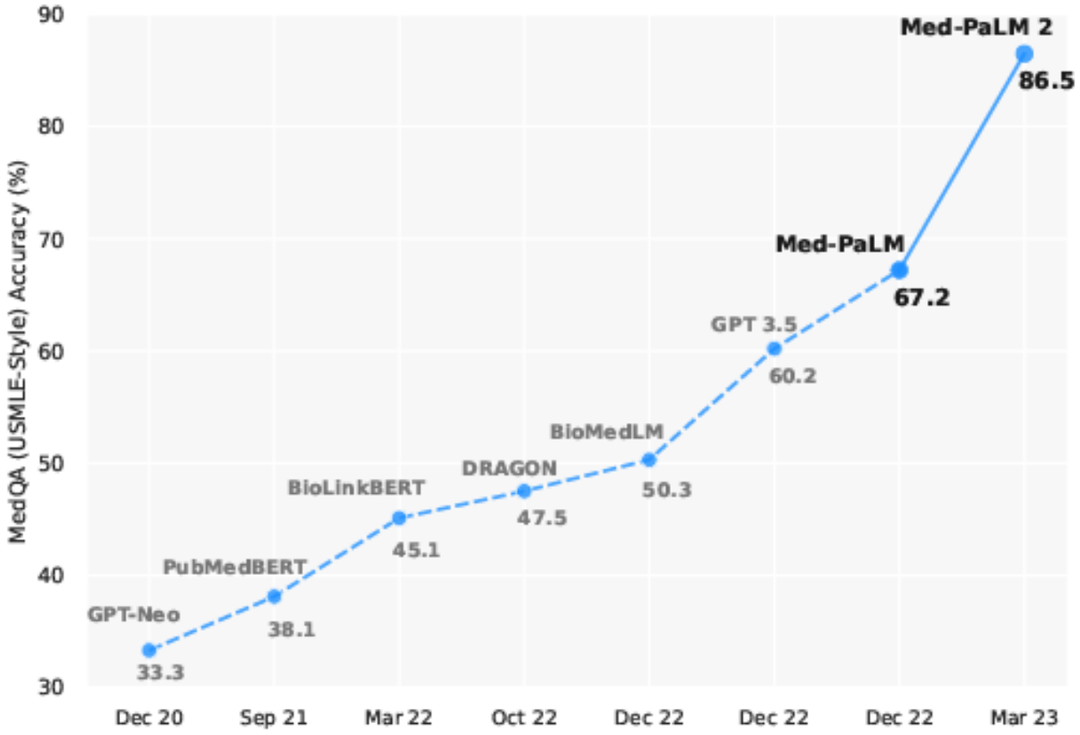

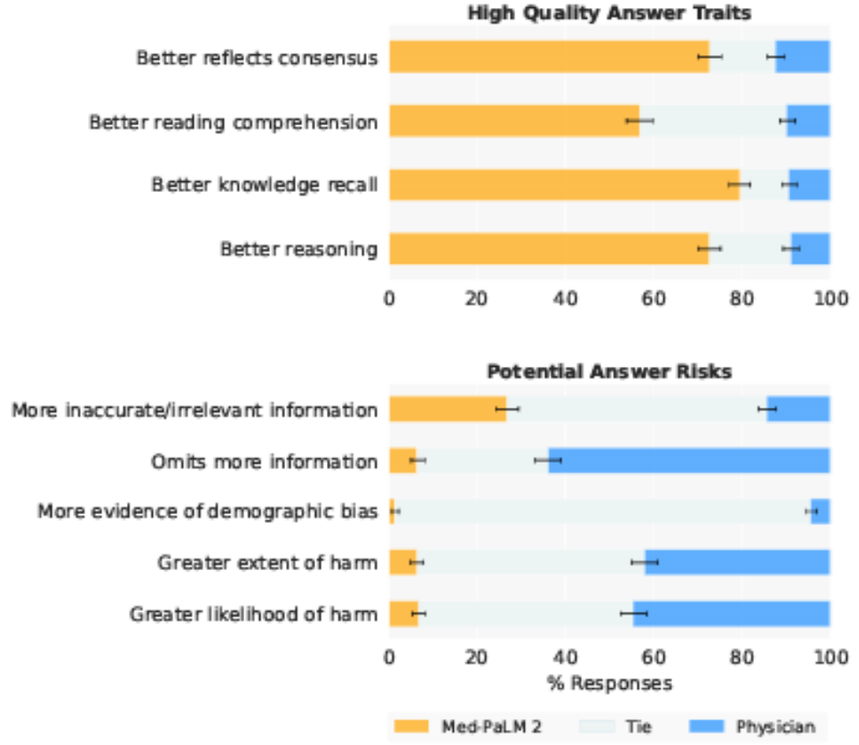

Source: Google. Toward Expert Level Medical Question Answering with LLMs.

manatt

Things AI Is Good At

AI Capability	Description	Example Clinical Scenario	Potential Role of AI
Identification	Identifying objects, patterns, and/or characteristics within data (often images).	A physician orders an X-ray for a patient who presents with pain, swelling, and limited leg mobility.	An AI tool reviews the X-ray and identifies an incidental nodule for further analysis by a radiologist.
Translation	Translating data inputs into another data type or data format (often between modalities or languages), often using natural language processing.	A radiologist reviews an MRI for a patient at-risk of breast cancer and dictates observations via an audio recording.	An AI tool converts the radiologist's audio dictation into a structured summary and applies the BIRADS classification scheme automatically. The AI produces a 'plain language' interpretation for the patient. The AI could also translate the report into a different language.
Summarization	Summarizing data inputs into shorter and more accessible outputs.	A patient is admitted to an emergency room in status epilepticus. A team of admitting healthcare providers review the patient's medical file to understand the patient's medical history, current medications, previous allergic reactions, and potential triggering factors.	An AI tool reviews the patient's medical history in totality, near- instantly identifying and summarizing key information for current clinical needs, such as recent medication changes affecting seizure threshold and a list of contraindicated drugs based on allergy history.
Prediction	Predicting or forecasting future events based on historical data and patterns.	A patient is discharged after hospitalization for heart failure.	Using historic heart failure readmission rates and the patient's clinical data, an AI tool predicts the risk of the patient's hospital readmission.
Suggestion	Providing recommendations, guidance, or advice. In some systems, suggestions may automatically lead to a specific downstream action.	A patient sees a provider every few months for a routine check-in; provider team conducts retrospective analysis of blood glucose measures from past few months.	An AI tool continually monitors a patient's blood glucose levels and (1) sends an alert to patient and provider when deviations occur and (2) provides recommended course of action (e.g., insulin level recommendation).

Source: American Medical Association. Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care. 2024.

Specialties: Current & Future AI Use Cases (Examples)

Specialty	Example AI use cases in practice today	Potential use cases in the future
Heart Care	• Detect arrhythmias, ischemia, and other heart abnormalities through electrocardiogram (ECG) analysis • Predict fractional flow reserve (FFR) from computed tomography (CT) images to aid in assessment of coronary artery disease • Use guidance during echo image analysis to optimize views and enable upskilling of practitioners • Triage data from remote patient monitoring devices (e.g., wearable sensors) and electronic health record (HER) systems to identify patients at highest risk for disease progression	• Triage data from imaging and EHR for early diagnosis and referral • Utilize ECG to assess progression of valvular disease • Assess personalized risk factors most strongly influencing cardiovascular outcomes for modification
Brain Care	• Analyze electroencephalography (EEG) studies to identify neurophysiological abnormalities and/or define the origin of seizures in the brain. • Detect and locate cerebrovascular abnormalities (e.g., ischemia) through continuous EEG analysis. • Analyze sleep polysomnography (PSG) to categorize sleep stages. • Identify stroke risk through the analysis of imaging data.	• Analyze multi-modal wearable data to predict, detect, and classify epileptic seizures • Analyze multiple sources of data to identify early biomarkers of neurodegenerative diseases (e.g., Alzheimer's, Parkinson's) • Triage data from remote patient monitoring devices (e.g., wearable sensors) to identify patients at the highest risk for brain injuries and post-brain injury complications
Eye Care	• Screen for and identify diabetic retinopathy through analysis of retinal imaging • Detect glaucoma through analysis of visual field tests and optical coherence tomography (OCT) scans • Quantify fluid on optical coherence tomography scans	• Autonomous screening and assistive diagnosis of all common forms of blindness through fundus images or OCT scans. • Predict cardiovascular and neurological risk factors through analysis of retinal fundus photographs
Pediatrics	• Monitor patient vitals to detect deviations and anticipate onset of critical illnesses (e.g., sepsis) • Analyze patient's medical history and other inputs such as vitals, images, videos, etc., to predict disease (e.g., autism spectrum disorder)	• Detect developmental delays in infants and children through analysis of video, speech, and/or written data • Predict risk of acute illness (e.g., pneumonia) earlier than current practice through X-ray analysis

Source: American Medical Association. Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care. 2024.

Known Issues with AI

Bias	Prejudices or unconscious biases in training datasets may inform AI model's outputs.
Explainability *Ability to explain how an AI output was generated from inputs.*	Newer models (e.g., foundational models) have low degrees of explainability due to large size and significant complexity.
Transparency *Ability to access information about an AI model's training data and model details*	Difficult for increasingly complex and evolving models to meet high standards of transparency, which are important for end-users in determining whether the model will work as expected.
"Hallucination"/ Confabulation	Generative AI models producing outputs that are either nonsensical or appear credible but are factually inaccurate.
Liability	Where liability falls (e.g., developer, organization, provider) is unsettled.
Coding & Payment	Nascent, but growing. Until recently, there was no common terminology to describe health care services or procedures delivered via AI, and no standardized way to pay for those tools and services.
Privacy & Security	AI development and training rely on access to large data sets; today, few, if any, technical controls are available to help end users specify how systems are trains and/or how data entered into AI systems are used and reused.
Regulation & Oversight	While the regulatory landscape is seeing significant uptick in activity at both the federal and state levels, significant gaps exist. Multiple parties (federal, state, organizational) will need to collaborate to ensure successful oversight.

Source: American Medical Association. Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care. 2024.

NIST AI Risk Management Framework

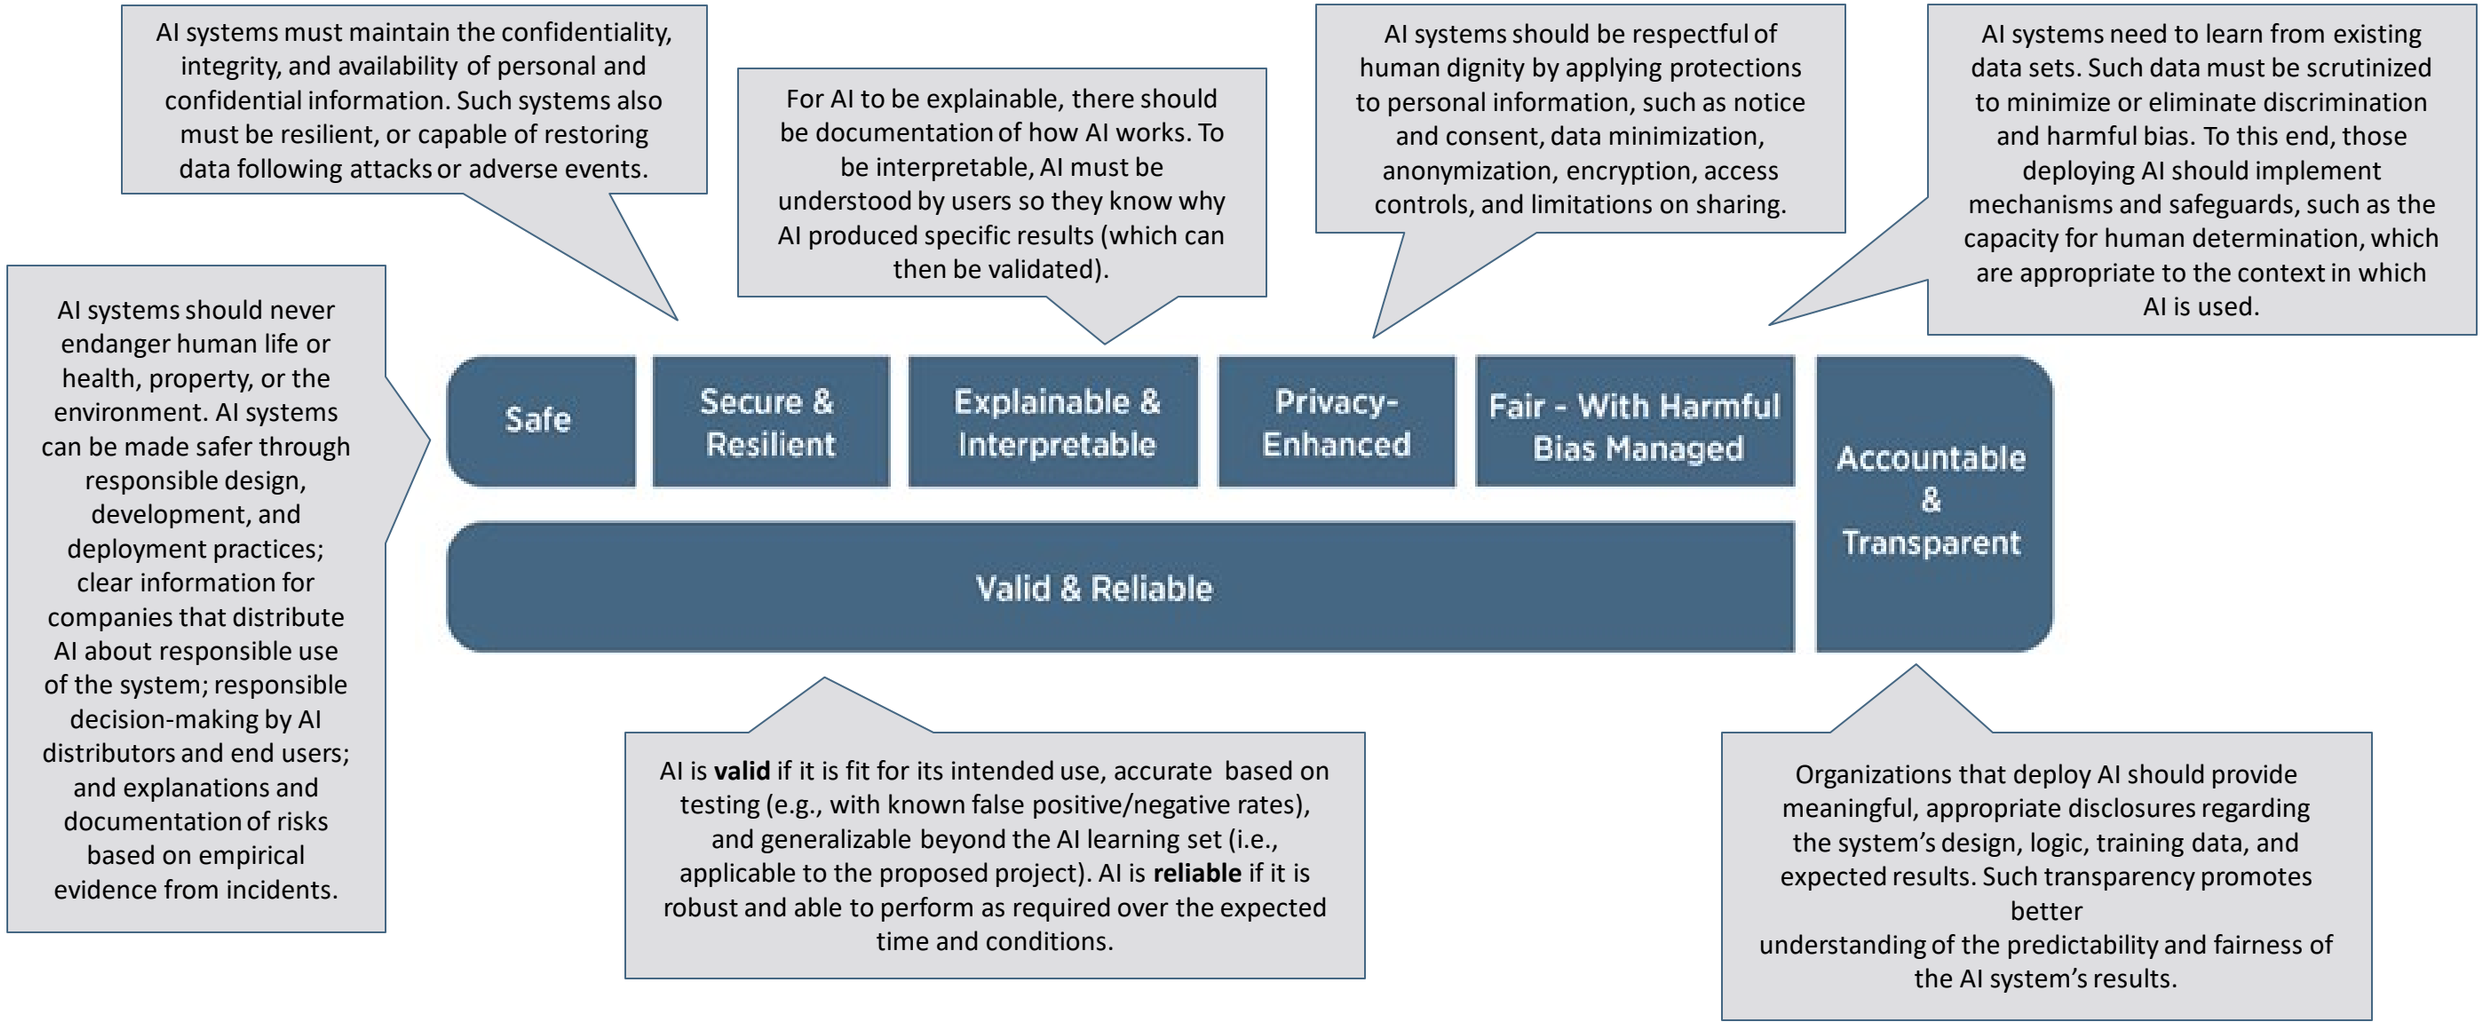

Source: U.S. National Institute of Standards and Technology (NIST). AI Risk Management Framework (RMF) 1.0. January 2023. (Adapted).

manatt

Health AI – Legal and Policy Issues

manatt

Legal and Regulatory Framework Governing AI

There is currently no federal law governing AI. In the 2023 legislative session, at least 25 states, Puerto Rico and the District of Columbia introduced artificial intelligence bills, and 15 states and Puerto Rico adopted resolutions or enacted legislation. Most state laws are aimed at studying AI (before regulating it).

There are many existing federal and state laws, regulations and agencies that do govern AI and company's use of AI.

- **Copyright and Trademark Laws**
- **HIPAA**
- **FTC**
- **Employment Laws**
- **Office of the National Coordinator**
- **State Privacy Laws**
- **Anti-Discrimination Laws**
- **Laws Governing Professional Conduct**
- **Laws Governing Licensed or Registered Activities**
- **Office of Civil Rights**
- **FDA**

Source: https://www.ncsl.org/technology-and-communication/artificial-intelligence-2023-legislation ; https://www.manatt.com/insights/newsletters/privacy-and-data-security/regulation-of-ai-systems-is-already-here-look-to.

Summary of Federal AI Activity To-Date

White House	AI Executive Order, AI Blueprint for an AI Bill of Rights, etc.
HHS Office of Civil Rights	Proposed Non-Discrimination Rule under 1557 of the Affordable Care Act
Centers for Medicare and Medicaid Services	Regulatory guidance for Medicare Advantage plans regarding use of clinical algorithms for medical necessity determinations
HHS Office of the National Coordinator	Transparency and risk management requirements for AI developers (HTI-1 Rule)
Food and Drug Administration	Guidance on AI-driven clinical decision support tools that should be regulated as medical devices
Department of Justice	Pending litigation over alleged use of AI to deny Medicare Advantage claims

Legal and Regulatory Framework Governing Health AI: Federal

	Role	*Anticipated Health AI Implications*	*Activities To-Date*	*Ongoing & Expected Activity*
Office of the National Coordinator for HIT	Leads national efforts to implement the use of certified HIT and the exchange of health information	Establish governance and transparency rules for how AI can be used in certified HITs (e.g., clinical decision tools, information blocking compliance)	HTI-1 Rule	HTI-2 Rule
Office of Civil Rights (OCR)	Enforces compliance with civil rights laws broadly, beyond health tech and AI	Implement rules that protect patients from discriminatory actions within the scope of medicine broadly, including use of AI tools (e.g., racial bias in use of photo-based AI clinical diagnosis tools)	Proposed 1557 Rule	Rule expected to be finalized Spring 2024
FDA	Regulates the development of various products including medical devices and prescription drugs	Regulate AI tools and software that meets the statutory definition of "medical device"	Non-binding guidance on clinical decision support software	Address gaps in Clinical Decision Support Software regulatory framework
CMS	Administers the Medicare program and partners with state governments to administer Medicaid, the Children's Health Program (CHIP), and health insurance portability standards	Propose regulations or issue guidance on how AI can be used within Medicare and/or Medicaid (e.g., conditions of payment, prior authorization). Establishes coverage and payment for AI solutions through various fee schedules	Regulatory guidance on MA medical necessity determinations	Issue further guidance or regulations on use of AI to administer Medicare/Medicaid
FTC	Enforces civil antitrust law and promotes consumer protection	Investigates privacy and security risks associated with AI and online tracking technologies that collect patient information (e.g., Meta/Facebook pixel, Google Analytics)	Letters to health systems/hospitals on risks of online tracking technologies	Continued attention on privacy and security in era of AI

Legal and Regulatory Framework Governing Health AI: States

	Purpose	*Anticipated AI Implications*
State Privacy Laws	State-specific laws that can place more protections on provider use of patient data.	Amend laws to address digital health and data-related gaps not covered by HIPAA (e.g., patient consent, data-sharing transparency requirements, governance structures for discrimination, etc.).
Laws Governing Licensed or Registered Activities	State-specific laws that dictate requirements to obtain and maintain various types of professional licenses.	Issue facility-based requirements for how AI can be used by licensed hospitals, physician practices, etc.
Laws Governing Professional Conduct	State-specific laws (licensure, scope of practice) that regulate how licensed medical professionals can practice and provide care.	State professional boards (Medicine, Nursing, etc.) may issue guidelines or statements regarding use of AI in clinical practice, ranging from expansive (i.e., the Board will not regulate use of AI, will rely on the clinician's discretion) to more limiting (i.e., it is professional misconduct to rely solely on AI for clinical decisions).

What about HIPAA?

- HIPAA provides the floor of privacy protection for patients' identifiable health information. State privacy laws build upon HIPAA.
- The rapidly evolving nature of health AI tools introduces new challenges in maintaining HIPAA compliance – particularly for AI tools that handle protected patient health information.
- Under HIPAA, any AI tool (particularly 3rd party tools) that use protected patient health information are required to execute a HIPAA Business Associate Agreement.
- Healthcare organizations should work with AI developers to understand how their AI tools work and ensure they are maintaining HIPAA compliance.

Pres. Biden Signed a Wide-Ranging Executive Order on AI in Oct 2023 Aimed at Promoting Responsible AI Innovation

OCTOBER 30, 2023

FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence

BRIEFING ROOM › STATEMENTS AND RELEASES

- **The EO aims to promote safety, privacy, equity, and innovation while ensuring responsible government use of AI.**
- **It lacks strong enforcement and instead directs various federal agencies to take actions that mitigate the potential harms of AI.**
- **More comprehensive AI regulation would require congressional action.**

Health Sector Implications

Much of the EO relates to cybersecurity and AI use in government and non-health sectors (see detailed appendix slides), however there are implications for the health care field:

- **Responsible Deployment of AI in Healthcare**. HHS will establish an AI Task Force to develop policies to responsibly deploy AI in the healthcare sector, including research and discovery, drug and device safety, healthcare delivery and financing, and public health. Outputs from the Task Force will undoubtedly impact a broad range of stakeholders, including providers, payers, device manufacturers, pharmaceutical companies, public health officials, regulators and patients.
- **Enforcement of Existing Laws to Prevent Harmful Use of AI**. Federal agencies are directed to enforce existing laws to mitigate against AI practices that result in unfair or deceptive business practices, privacy violations, or discrimination. This will apply to healthcare providers and insurers that use AI in relation to serving their patients and members.
- **Investment in AI Research.** Federal agencies are tasked with fostering AI innovation, including offering research grants to study AI in healthcare. Healthcare researchers and academic medical centers will likely see more research funding, computational resources, data, and training opportunities.
- **Identification of AI-Related Errors in Healthcare Delivery.** HHS and the Departments of Veteran Affairs and Defense must create a framework to identify and log clinical errors resulting from AI in healthcare settings. This effort may help lay the groundwork for identifying and preventing AI errors across all providers.

In December 2023, 28 providers and payers voluntarily committed to the EO's principles for use in healthcare

Sources: AI Executive Order, Manatt Newsletter on Executive Order

manatt

HHS/ONC Finalized "HTI-1" Rule Imposing AI Transparency and Risk Management Requirements on Certified HIT Developers

In December 2023, HHS, through the Office of the National Coordinator for Health Information Technology (ONC), finalized its "Health Data, Technology and Interoperability (HTI-1)" rule, which includes new requirements for AI and other predictive algorithms embedded within EHR systems and used by healthcare providers.

- **Purpose**: Promote the development of predictive algorithms that are fair, appropriate, valid, effective and safe (FAVES) and ensure that the AI used by health care providers can be trusted.

- **New Requirements for Developers of ONC-certified HIT that use Predictive Decision Support Interventions (DSI)*** :

  1. Transparency. If certified HIT* uses predictive DSI, the HIT developer must make available to the software users (i.e., providers) detailed information about the predictive DSI, including:
     - The purpose of the intervention;
     - Funding sources for the intervention's development;
     - Exclusion and inclusion criteria that influenced the training data set;
     - The process used to ensure fairness in development of the intervention; and
     - A description of the external validation process.

  2. Risk Management**.** Predictive DSI must be subject to: 1) an analysis of potential risks and adverse impacts associated with its validity, reliability, robustness, fairness, intelligibility, safety, security and privacy; 2) practices to mitigate risks; and, 3) policies and implemented controls for governance (including how data are acquired, managed and used).

These requirements aim to ensure that health systems, clinicians and other users of the predictive DSI understand the software that is being made available to them and that it is aligned with other government actions aimed at increasing transparency with regard to AI tools.

*Note: The rule only pertains to developers of ONC-certified health information technologies. Predictive algorithms used by providers that are not offered as part of certified HIT are outside the regulation's scope. Predictive DSI is "technology that supports decision-making based on algorithms or models that derive relationships from training data and then produces an output that results in prediction, classification, recommendation, evaluation, or analysis (i.e., technologies that employ AI).

Source: HHS Final Rule, Manatt Health Newsletter

HHS is expected to propose "HTI-2" rule later this year

manatt

“HTI-1” Rule: Implications for Providers

The ONC’s rule is the first health care-focused regulation targeting providers’ use of AI, yet it primarily impacts developers.

- The new rule impacts certified HIT developers more than health care providers, who can use unregulated AI tools outside of certified HIT.
- Health care providers should closely track predictive DSI requirements being imposed by the ONC; they may foreshadow how HHS may regulate health care providers’ use of AI through other mechanisms (e.g., Medicare Conditions of Participation).
- Additionally, HHS’ approach to regulating providers’ use of AI may serve as a blueprint for states.

Other Federal Regulations Governing Use of AI in Health

- **Proposed regulations under Section 1557 of the Affordable Care Act - Non-Discrimination**: Proposed new regulation at 45 CFR 92.210 "A covered entity must not discriminate on the basis of race, color, national origin, sex, age, or disability in its health programs and activities through the use of clinical algorithms in its decision-making". Nondiscrimination in Health Programs and Activities, 87 FR 47824.

Notable Implications for Providers

  - In this proposed rule, HHS notes that *"[w]hile covered entities are not liable for clinical algorithms that they did not develop, they may be held liable under this provision for their decisions made in reliance on clinical algorithms."*
  - Health care providers should ensure they understand how AI is being used within their organization, how the AI tools were developed, and demonstrate that their AI has been tested for bias and discrimination.

Manatt Insights:

- Final rule expected by April 2024
- OCR's intent is to make sure clinicians use algorithms thoughtfully and in a way that is not discriminatory
- Clinicians should be aware of algorithmic biases within AI tools used in their practice (e.g., kidney function algorithm that accounts for race)
- Enforcement will be challenging

- CMS provided regulatory guidance, stating that **Medicare Advantage Plans** "**must** ensure that they are **making medical necessity determinations** based on the circumstances of the **specific individual**, as outlined at § 422.101(c), as **opposed to using an algorithm or software that doesn't account for an individual's circumstances**." 88 FR 22120, 22195

FDA Regulation of Clinical Decision Support (CDS) Software

21st Century Cures Act excluded certain medical software, including CDS, from definition of "medical device" in the Federal Food, Drug, and Cosmetic Act (FFDCA) (i.e., FDA lacks jurisdiction over these products).

The following technologies are statutorily excluded from the "medical device" definition:

Software intended for administrative support of a health care facility, such as processing bills and claims, scheduling, inventory management, analysis of historical data to predict utilization, and determination of benefit eligibility.

Software intended for maintaining or encouraging a healthy lifestyle where the use is unrelated to the diagnosis, cure, mitigation, prevention, or treatment of a disease or condition (i.e., general health and wellness).

Software intended to serve as electronic patient records, if such records are created and used by health care providers and constitute health information technology as certified under Public Health Service Act § 3001(c)(5), and the software does not interpret or analyze patient records for the purpose of diagnosis, cure, mitigation, prevention, or treatment of a disease or condition.

Software intended to transfer, store, convert formats, or display (but not analyze) laboratory and device data, results, associated findings by a health care professional, and general background information about a test or device.

Clinical decision support ("CDS") software that displays/analyzes medical information and makes recommendations for health care professionals regarding prevention, diagnosis, or treatment of a disease/condition.

However, such software:

- Cannot be processing or analyzing medical images (e.g., CT scans, MRIs, X-rays), data from *in vitro* diagnostic devices, or "a pattern or signal from a signal acquisition system"; and
- Must be transparent and its recommendations reviewable, such that it does not serve as the sole basis for a health care professional's determination regarding a particular patient.

Source: Section 520(o), Federal Food, Drug, and Cosmetic Act; 21st Century Cures Act (P.L. 114-255), Section 3060(a).

The FDA Has Issued Non-Binding AI-Related Guidance and Proposed Frameworks

This Action Plan outlines 5 steps the FDA intends to take to further oversight for AI/ML-based SaMD, including:

1. Further developing the proposed regulatory framework, including through issuance of draft guidance on a predetermined change control plan (for software's learning over time);
2. Supporting the development of good machine learning practices to evaluate and improve machine learning algorithms;
3. Fostering a patient-centered approach, including device transparency to users;
4. Developing methods to evaluate and improve machine learning algorithms; and
5. Advancing real-world performance monitoring pilots.

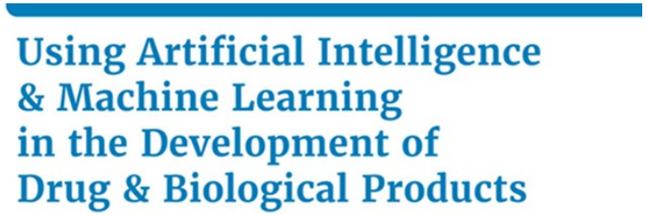

This discussion paper provides an overview of current and potential future uses for AI/ML in the development of drug and biological products, and addresses concerns and risks associated with these innovations.

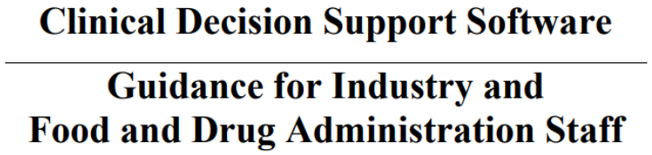

This FDA guidance on AI-driven clinical decisions support tools recommends that some AI-powered tools, like sepsis prediction devises, should be regulated as medical devices.

Explored in more detail on subsequent slides

The FDA also maintains a details list of AI/ML-Enabled Medical Devices, available here.
As of October 19, 2023, over 700 devices with embedded AI/ML have been authorized.

How does FDA CDS Software Guidance Impact Generative AI?

Overview of FDA Approach to Healthcare AI

- Numerous types of generative AI seek to both advance clinical care and alleviate provider burden
  - Diagnostic tools, Risk Prediction, Triage Tools, etc.
- As AI technology evolves, FDA is striving to evolve its regulatory approach within the confines of the statute (i.e., AI/ML-Based SaMD Action Plan)
- **FDA may only regulate software that meets the statutory definition of "medical device"**
  - **AI-Generated clinical summaries that do not provide healthcare "recommendations" are likely to be considered exempt as software "intended to serve as electronic patient records";**
  - **If AI-generated summaries provide recommendations, yet are intended to be "Augmented Intelligence," (i.e., not intended to be sole basis of provider decision-making), then likely to be non-device CDS software**

Case Study: AI-Generated Clinical Summaries

- AI-driven clinical support tools that utilize large language models (LLMs) to summarize clinical notes, medications, and other forms of patient data
  - Range from LLM-generated summaries from audio-recorded patient encounters to future state summaries of patient information across EHRs (clinical "snapshots")
- Many of the LLMs that summarize clinical notes, medications and other patient info may fall outside the definition of devices that require FDA oversight, per the new guidance
- LLM-generated summaries are highly variable based on inputs and can include errors
- Provider reliance on summaries that have varying degrees of accuracy and detail could influence subsequent decisions on patient care

Key Takeaway: There are notable gaps in how the FDA defines medical devices relative to new AI tools - LLMs could impact clinical decision-making, yet will not have FDA oversight

In January 2024, the GAO published a report highlighting necessary enhancements to the FDA's authority to oversee AI/ML devices.
The GAO urged the FDA to specify legislative changes that would help address potential regulatory gaps in FDA-oversight of emerging technologies.

FDA Guidance on Clinical Decision Support (CDS) Software

September 2022 FDA guidance provided additional detail on how FDA interprets 21st Century Cures' provision on CDS Software and includes numerous examples of non-device and device CDS software functions.

To be considered CDS Software (and therefore exempt from FDA definition of "medical device"), software functions meet all of the following four criteria:

1. Not intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system;

2. Intended for the purpose of displaying, analyzing, or printing medical information about a patient or other medical information (such as peer-reviewed clinical studies and clinical practice guidelines);

3. Intended for the purpose of supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition; and

4. Intended for the purpose of enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.

Source: FDA CDS Guidance (Sept 2022)

Key Takeaways: Gaps in the Oversight and Regulation of Clinical AI Tools

- **Gap in Oversight of Certain AI Tools due to Statutory Definitions.** The FDA does not have regulatory oversight for certain uses of AI/ML clinical tools.
- **Future of Regulatory Oversight for AI Remains Uncertain.** The FDA is not adequately resourced to provide sufficient oversight of currently exempt AI/ML tools. States are similarly not well-positioned or resourced to be regulating; a state-based approach would result in patchwork AI landscape.
- **GAO Directed the FDA to Provide Input to Congress on How to Address Gap**. Congress would need to take action (i.e., expand the definition of "medical devices" within statute) for the FDA to oversee the development of currently excluded AI/ML clinical tools. Recent GAO report suggested the FDA specify legislative actions to address current regulatory gap.
- **Some Stakeholders are Seeking to Address the Need**:
  - The Coalition for Health AI (including the FDA and ONC) proposed establishing AI Assurance Labs to evaluate AI models based on an agreed-upon set of principles and best practices.
  - Some states are introducing legislation that would require states to adopt guidelines for the development/deployment of AI tools (e.g., registering models with the state).

Implications for FQHCs

manatt

Conclusions

- AI in health care is making rapid progress and is likely here to stay
- AI will never replace clinicians — but clinicians who use AI will be better positioned to care for their patients than those who don’t
- AI model performance is improving at a rate that will soon match/surpass human capabilities for many tasks
- AI requires ongoing oversight and continuous validation after deployment
- Robust governance is essential as regulators catch up → adopting the CHCANYS Model AI Policy is a good first step
- Collaboration between technologists and clinicians is essential to effective AI implementation in practice