MAFP: AI IN FAMILY MEDICINE

ENHANCING CARE, EFFICIENCY, AND CONNECTION

Jason Greenwood MD MS
Marc Matthews MD
Adria Hoffman PhD
David Rushlow MD

OUR TEAM

Jason Greenwood MD MS
Assistant Professor - FM
Director of Clinical Informatics and AI – FM RST

Marc Matthews MD
Assistant Professor – FM
Director of Innovation – FM RST

Adria Hoffman, PhD
Sr. Education Specialist

David Rushlow MD
Assistant Professor – FM
Chair – FM Rochester MN

AI is Everything?

GUIDING QUESTIONS

- What are some of the technologies referred to as AI that are useful in our work as physicians?
- How are AI tools developed?
- How will I determine whether and how to use an AI tool in the clinical setting?

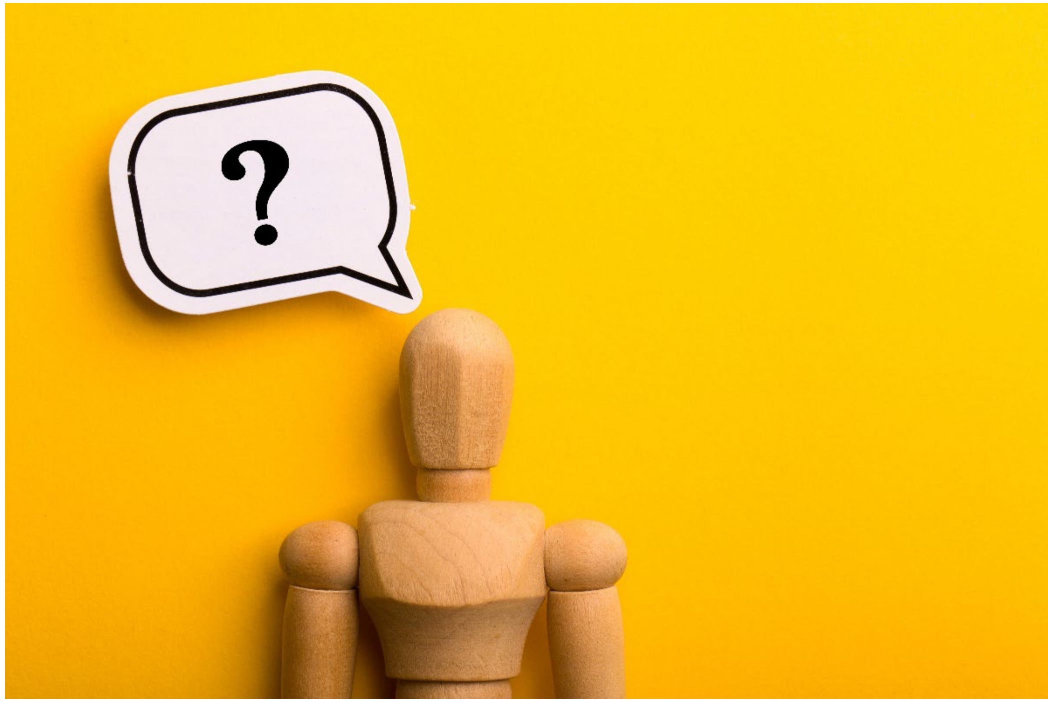

LEARNING OBJECTIVES

- Define key AI terminology
- Define the rationale for using AI
  - Identify questions to determine an AI tool's affordances and constraints.
  - Apply these questions to critically evaluate whether or not a particular tool is beneficial to you or your patients
- Describe how various AI technologies are developed
- Identify opportunities for AI to decrease burden
- Describe AI model card usage

THE "BIOCHEMISTRY" OF AI

THE COMPUTER/SERVER

Hard drive – Data storage
Like the file cabinet of records
The bigger the cabinet the longer back we can store

Random Access Memory (RAM) - Data Workspace
The "desk" or workspace of the program/model
The larger the desk the more we can work on at once
and the larger projects we can take on

The Central or Graphics Processing Unit (CPU/GPU) - Brain
The "people" working on the problem at the "desk"
The more specialized people we have working on the problem,
the faster we can accomplish the goal

ALGORITHMS
NOT NECESSARILY "AI"

- "a process or set of rules to be followed in calculations or other problem-solving operations, especially by a computer" (Oxford)

- They work using **basic mathematical or If-then-else** operations of a basic CPU

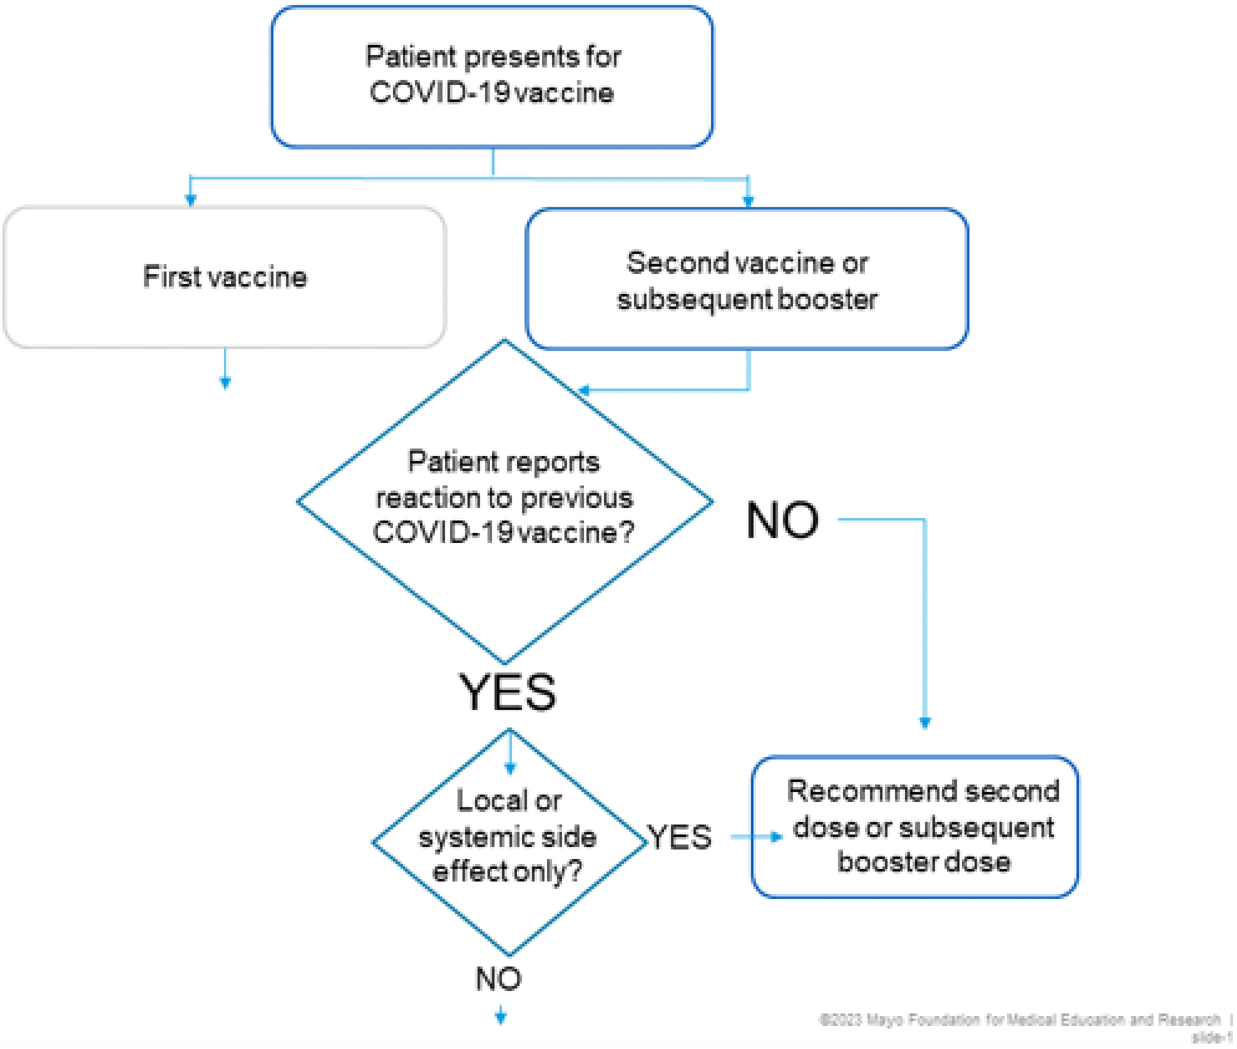

WHAT IS AI?

- Artificial Intelligence: **An umbrella term for many kinds of technologies**, including some that don't currently exist but receive a great deal of media attention.

- How to think about AI: Computer-based systems and tools that are designed by humans to **efficiently find patterns in large data sets** and to **automate tasks** traditionally accomplished manually.

A program able to perform tasks that normally require human intelligence

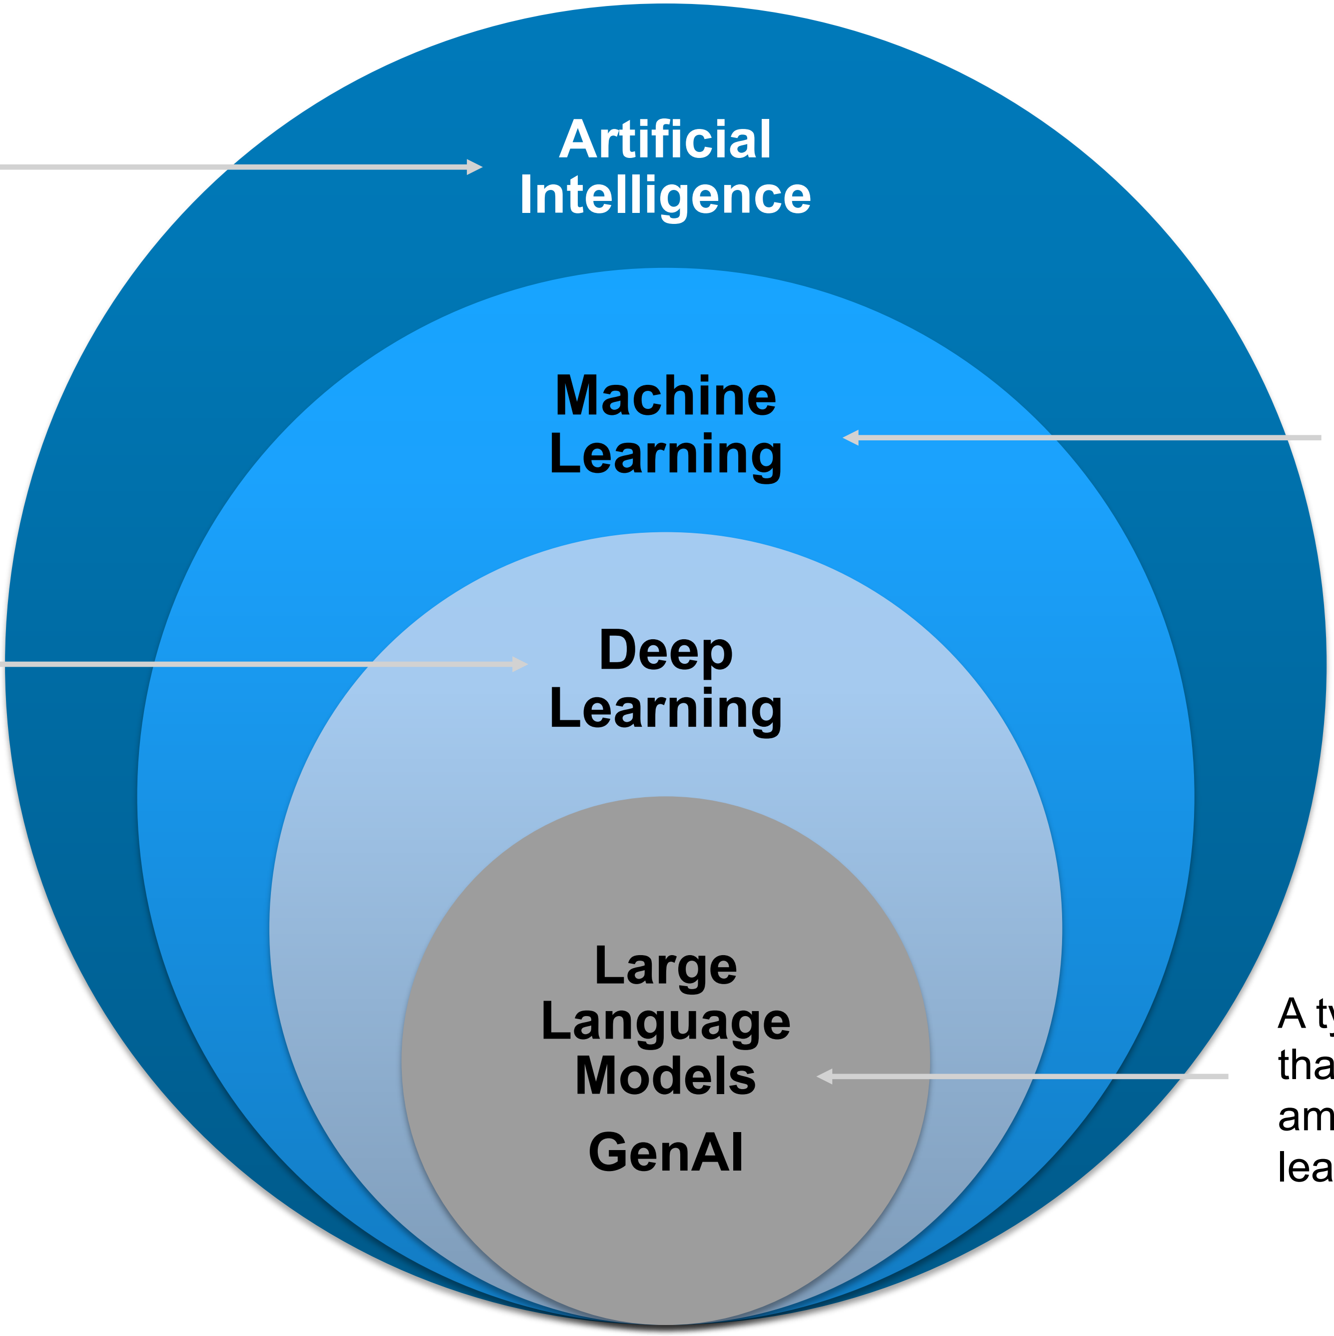

Able to find patterns in simpler sets of data (can be more explainable)

A subset of machine learning where the program can find more complex patterns based on larger sets of data (Black Box)

A type of deep learning model that has been trained on large amounts of data and is able learned complex patterns

LARGE LANGUAGE MODELS

- **Large Language Models (LLM)**
  - Examples: GPT-4o, llama 3, Gemini 2.0
  - How & what: predict the probable next word based on existing narrative data
- **ChatGPT, Gemini, Copilot, etc.**
  - What: public interface for GPT based models from Open AI
  - Do not use this for protected health information (PHI) or confidential business information (CBI) (unless your organization has a business agreement (BA) that protects data from being shared outside of your organization)
- **Prompts**
  - Text/data input/instruct an LLM to get desired output
- **Prompt Engineering**
  - Generative AI output quality is highly variable. Prompt engineering carefully crafts inputs for optimized outputs.

 |

AP WORLD U.S. POLITICS SPORTS ENTERTAINMENT BUSINESS SCIENCE FACT CHECK ODDITIES BE WELL NEWSLE

TECHNOLOGY

Researchers say an AI-powered transcription tool used in hospitals invents things no one ever said

Groundbreaking BBC research shows issues with over half the answers from Artificial Intelligence (AI) assistants

Conducted over a month, the study saw the BBC test four prominent, publicly available AI assistants

Published: 11 February 2025

IMAGING INFORMATICS | ARTIFICIAL INTELLIGENCE

ChatGPT perpetuates racial and gender biases

Will Morton

Jan 2, 2024

AREAS AI INFLUENCING

- Predictive/Identifying
  - Imaging
    - Radiology, Pathology, Dermatology, etc.
  - Signal
    - ECG, EEG, Voice, PFT, etc.
  - Discrete
    - Labs, Surveys, Drug development, etc.
- Generative
  - Summarization
    - Ambient, Patient History, etc.
  - Information Retrieval
    - Literature, consultation, knowledge, etc.

THE "PHYSIOLOGY" OF AI

- **Why do we need AI?**
  - Healthcare is progressing at an exponential rate (literature, technology, treatments, etc.)
  - Humans can only process information in a linear fashion – we need help to keep up
  - Decrease burden to allow PCPs to operate at top of license and meet increasing populations needs

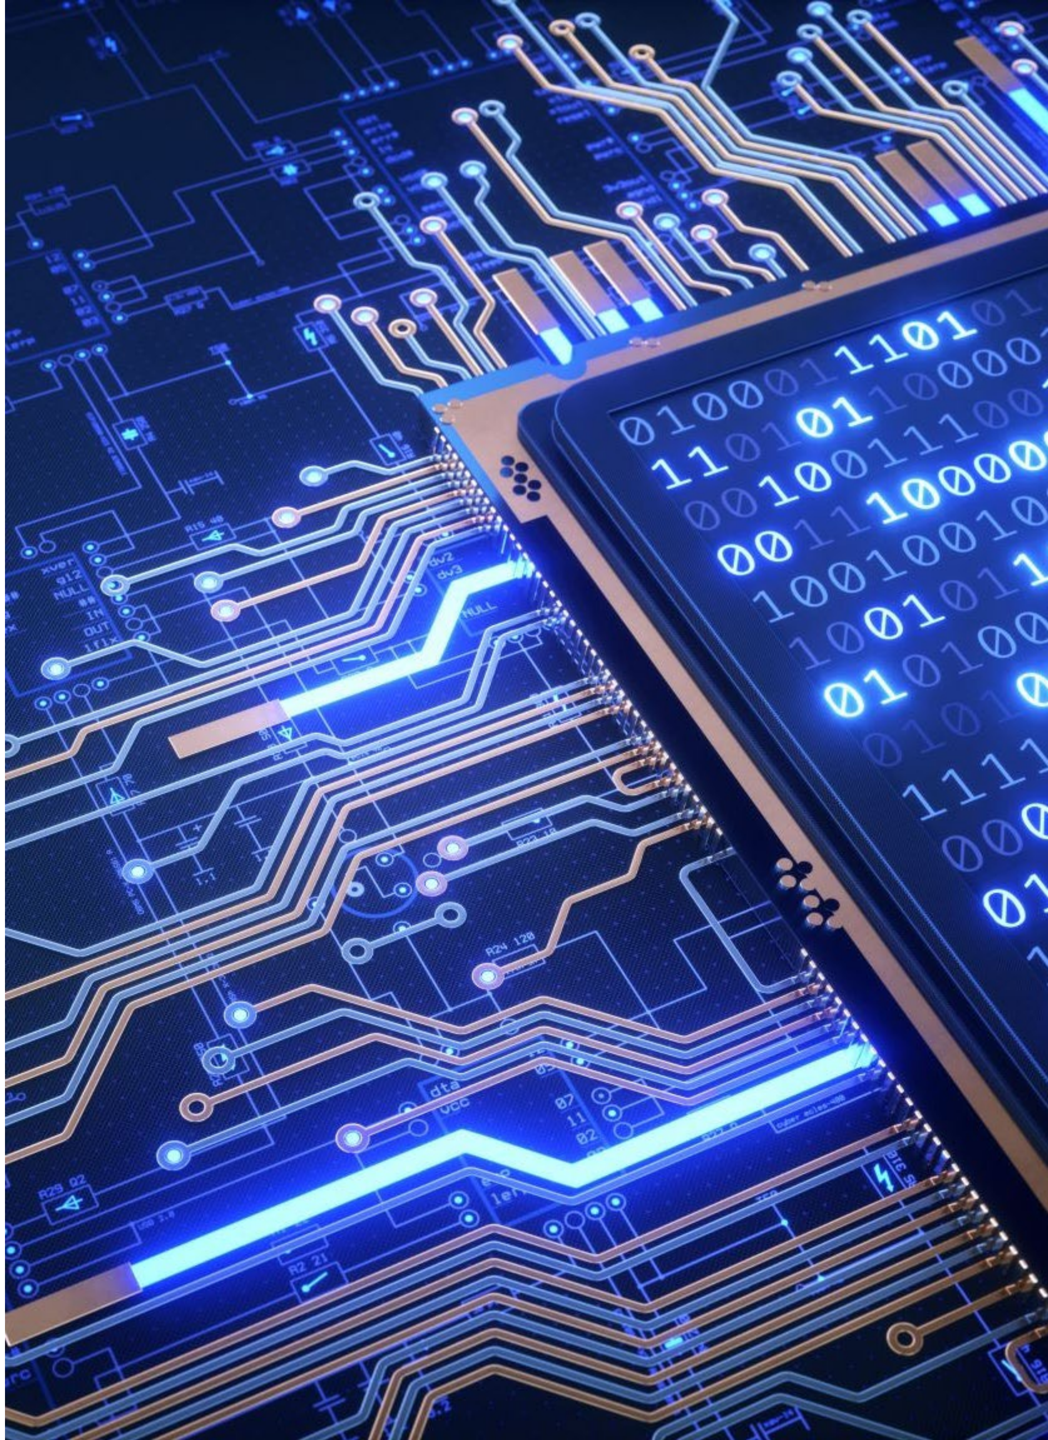

DEVELOPING AI MODELS - STEPS

- **Define the problem space/benefit**
- **Ready to deploy**
  - Ensure model is appropriate for your populations/clinic
  - Run silently in real time
  - Pilot deploy before full deploy
- **Post deploy monitoring**
  - Ensure it is performing as expected
- **Re-evaluate model yearly**
  - Practice, methods or dependent model updates

Because humans develop and train AI using existing data, these tools are just as “biased” as their human developers/providers. We must work to mitigate this bias.

 |

AI Model Card (Soup Can Label)

Model Facts **Model name:** Deep Sepsis **Locale:** Duke University Hosp

Approval Date: 09/22/2019 **Last Update:** 01/13/2020 **Version:** 1.0

Summary

This model uses EHR input data collected from a patient's current inpatient encounter to estimate the probability that the patient will meet sepsis criteria within the next 4 hours. It was developed in 2016-2019 by the Duke Institute for Health Innovation. The model was licensed to Cohere Med in July 2019.

Mechanism

- **Outcome** ....sepsis within the next 4 hours, see outcome definition in "Other Information"
- **Output** ....0% - 100% probability of sepsis occurring in the next 4 hours
- **Target population** ....all adult patients >18 y.o. presenting to DUH ED
- **Time of prediction** ....every hour of a patient's encounter
- **Input data source**....electronic health record (EHR)
- **Input data type** ....demographics, analytes, vitals, medication administrations
- **Training data location and time-period** ....DUH, diagnostic cohort, 10/2014 – 12/2015
- **Model type**.... Recurrent Neural Network

Validation and performance

	Prevalence	AUC	PPV @ Sensitivity of 60%	Sensitivity @ PPV of 20%	Cohort Type	Cohort URL / DOI
Local Retrospective	18.9%	0.88	0.14	0.50	Diagnostic	arxiv.org/abs/1708.05894
Local Temporal	6.4%	0.94	0.20	0.66	Diagnostic	jmir.org/preprint/15182
Local Prospective	TBD	TBD	TBD	TBD	TBD	TBD
External	TBD	TBD	TBD	TBD	TBD	TBD
Target Population	6.4%	0.94	0.20	0.66	Diagnostic	jmir.org/preprint/15182

Credit: Duke University

- Training Inputs (patient demographics, AI models, etc.)
- If it requires retraining
- Performance Data
- Last Updated/Reviewed
- 510K Approval Status (FDA)

Your Clinical Decision Support Software: Is It a Device?

FDA

The FDA issued a guidance, Clinical Decision Support Software, to describe the FDA's regulatory approach to Clinical Decision Support (CDS) software functions. This graphic gives a general and summary overview of the guidance and is for illustrative purposes only. Consult the guidance for the complete discussion and examples. Other software functions that are not listed may also be device software functions. *

Your software function must meet all four criteria to be Non-Device CDS.

Summary interpretation of CDS criteria

1. Your software function does **NOT** acquire, process, or analyze medical images, signals, or patterns.
2. Your software function displays, analyzes, or prints medical information normally communicated between health care professionals (HCPs).
3. Your software function provides recommendations (information/options) to a HCP rather than provide a specific output or directive.
4. Your software function provides the basis of the recommendations so that the HCP does not rely primarily on any recommendations to make a decision.

Your software function may be non-device CDS.

Non-Device Examples

Non-Device examples display, analyze, or print the following examples of medical information, which must also not be images, signals, or patterns:

- Information whose relevance to a clinical decision is well understood
- A single discrete test result that is clinically meaningful
- Report from imaging study

AND

Non-Device examples provide:

- Lists of preventive, diagnostic, or treatment options
- Clinical guidelines matched to patient-specific medical info
- Relevant reference information about a disease or condition

AND

Non-Device examples provide:

- Plain language descriptions of the software purpose, medical input, underlying algorithm
- Relevant patient-specific information and other knowns/unknowns for consideration

Device Examples

Device examples acquire, process, or analyze:

- Signal acquisition systems
- In vitro diagnostics
- Magnetic resonance imaging (MRI)
- Next Generation Sequencing (NGS)
- Continuous Glucose Monitoring (CGM)
- Computer aided detection/diagnosis (CADe/CADx)

OR

Device examples display, analyze or print:

- Continuous signals/patterns
- Medical images
- Waveforms (ECG)
- More continuous sampling (aka – a signal or pattern)

OR

Device examples provide:

- Risk scores for disease or condition
- Probability of disease or condition
- Time-critical outputs

OR

Device examples:

- Basis of recommendations is not provided

Your software function is a device.

***Disclaimer:** This graphic gives a general overview of Section IV of the guidance ("Interpretation of Criteria in Section 520(o)(1)(E) of the FD&C Act"). Consult the guidance for the complete discussion. The device examples identified in this graphic are illustrative only and are not an exhaustive list. Other software functions that are not listed may also be device software functions.

EXAMPLES OF AI
ADMINISTRATOR APPLICATIONS

- **Personnel/HR**
  - Hospital/clinical staffing needs predictions
- **Meetings/Email/Information**
  - Meeting summarization, email summarizations, email prioritization
  - Chatbot for documentation
- **Financial**
  - Predicted financial burdens/success internal and external, budgets requirements
  - Estimated patient growth

EXAMPLES OF AI
CLINICAL APPLICATIONS

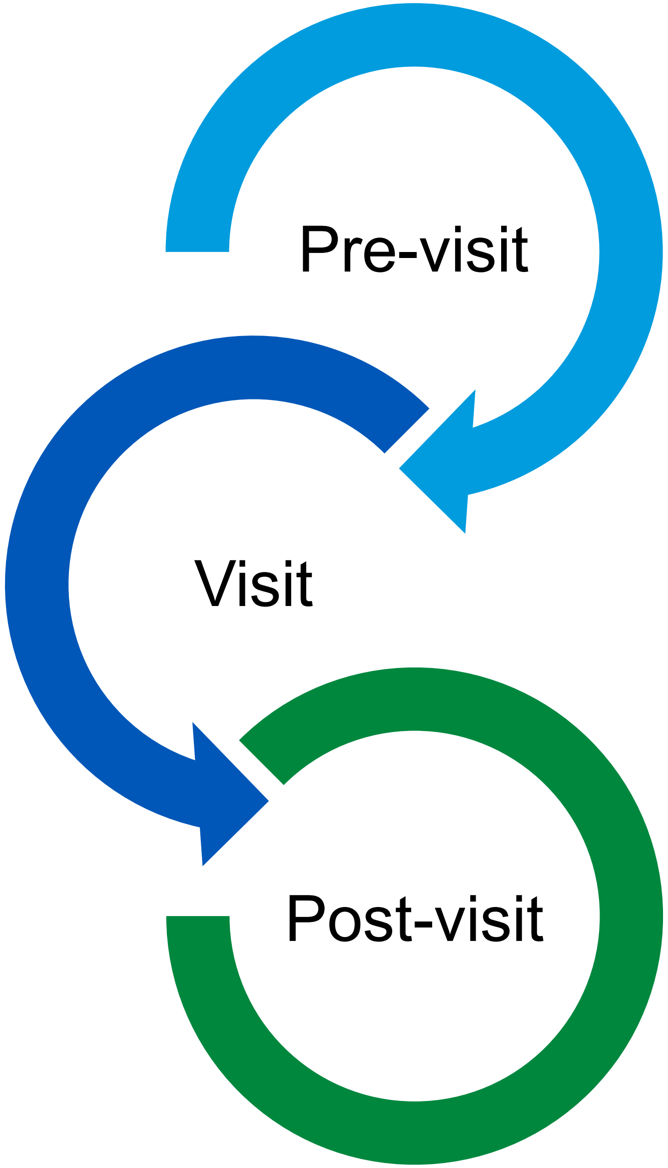

Pre-visit

- **Note/ Problem summarization (27% of provider time)**
- Predictive models for patient disease risk – help prioritize

Visit

- **Ambient documentation (24% of provider time)**
- HCC, Preventive care, Quality Measures

Post-visit

- **Augmented replies (23% of provider time)**
- Results suggested follow-up/ interventions – decrease referrals

Async-visit

- Augmented patient replies and predict need for return
- Clinical research summarization/interventions
- Care for more patients with less resources – augmenting care

AMBIENT DOCUMENTATION

- **Present**
  - Multiple vendors
    - Read the contracts!
  - Some are able to flow into the EHR
- **ROI**
  - Saving as much at 80% of documentation time
  - More visits closed same day
  - Decreasing burden
- **Upcoming**
  - Pend up orders, suggest problem list changes, problem-based documentation, and more

GENERATIVE REPLIES

- **Present**
  - Create a draft reply to a patient message in some cases with an empathetic and adjusted reading level response
- **ROI**
  - Can cut reply times by days/hours
  - Can cut message time by 30-60 secs
- **Upcoming EHR and Present Vendors**
  - Drafts for many other messages/ results, letters, denials, etc.

SUMMARIZATION

- **Present**
  - Able to create a summary of patient's hospital course, since your last visit, etc.
  - Be able to summarize literature
  - Multiple vendors and built into some EHRs
- **ROI**
  - Up to a 40% decrease prep time
  - Some as high as 98% felt better prepared
- **Upcoming EHR and Present Vendors**
  - Full chart summary, suggested follow up items, suggested HCC coding, integration with literature summarization, etc.

DATA EXTRACTION

- **Present**
  - Many vendors, homegrown tools and EHRs
  - Ensure to validate with your own data
- **ROI**
  - Able to help alleviate the chart diving burden
  - Can be used to trigger additional Clinical Decision Support (CDS)
  - Helpful in research
- **Upcoming EHR and Present Vendors**
  - Expect more fields, incidental findings, etc.

PREDICTIVE MODELS

- **Present**
  - Many institutions, vendors and EHRs offer solutions to predict outcomes/disease
  - Be sure to be at the table during development and implementation
- **ROI**
  - Able to reduce adverse outcomes
  - Could help predict resource needs
  - Might increase referrals/testing
- **Upcoming EHR and Present Vendors**
  - More and more models to come with the hope of more generalizability

Contents lists available at ScienceDirect
Data in brief
journal homepage: www.elsevier.com/locate/dib

Data Article

Clinical trial design data for electrocardiogram artificial intelligence-guided screening for low ejection fraction (EAGLE)

Xiaoxi Yao [a, b, c, *], Rozalina G. McCoy [a, d], Paul A. Friedman [c], Nilay D. Shah [a, b], Barbara A. Barry [b], Emma M. Behnken [e], Jonathan W. Inselman [a], Zachi I. Attia [c], Peter A. Noseworthy [c]

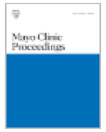

Mayo Clinic Proceedings
Volume 97, Issue 7, July 2022, Pages 1326-1336

Original article

Machine Learning Techniques Differentiate Alcohol-Associated Hepatitis From Acute Cholangitis in Patients With Systemic Inflammation and Elevated Liver Enzymes

Joseph C. Ahn MD [a], Yung-Kyun Noh PhD [a c], Puru Rattan MD [a], Seth Buryska BS [a], Tiffany Wu MD [a], Camille A. Kezer MD [b], Chansong Choi MD, MS [b], Shivaram Poigai Arunachalam PhD [a], Douglas A. Simonetto MD [a], Vijay H. Shah MD [a b], Patrick S. Kamath MD [a]

Volume 28, Issue 6
June 2021

Article Contents

JOURNAL ARTICLE

Using machine learning to improve the accuracy of patient deterioration predictions: Mayo Clinic Early Warning Score (MC-EWS)

Santiago Romero-Brufau, Daniel Whitford, Matthew G Johnson, Joel Hickman, Bruce W Morlan, Terry Therneau, James Naessens, Jeanne M Huddleston

Journal of the American Medical Informatics Association, Volume 28, Issue 6, June 2021, Pages 1207–1215, https://doi.org/10.1093/jamia/ocaa347
Published: 26 February 2021 **Article history** ▾

Mayo Clinic Proceedings
Volume 99, Issue 2, February 2024, Pages 260-270

Original article

Machine Learning for Diagnosis of Pulmonary Hypertension by Echocardiography

Vidhu Anand MBBS [a], Alexander D. Weston PhD [c d], Christopher G. Scott MS [b], Garvan C. Kane MD, PhD [a], Patricia A. Pellikka MD [a], Rickey E. Carter PhD [c d]

ARTICLES · Volume 4, Issue 9, E632-E645, September 2022 · *Open Access* Download Full Issue

Development of a multiomics model for identification of predictive biomarkers for COVID-19 severity: a retrospective cohort study

Seul Kee Byeon, PhD [a,†] · Anil K Madugundu, PhD [a,j,k,l,†] · Kishore Garapati, MBBS [a,j,k,l] · Madan Gopal Ramarajan, MD [a,j,k,l] · Mayank Saraswat, PhD [a,l] · Praveen Kumar-M, MD [m] · et al. Show more

DATA ANALYSIS

- **Present**
  - Vendors and EHRs offer their own solution
- **ROI**
  - Less of a learning curve with understanding the data
  - Gain further insights not previously seen in data
- **Upcoming EHR and Present Vendors**
  - Possibly anticipate outcomes based on current trends in data
  - Might cause increased worry or excitement

TAKE AWAYS

- Not everything is AI and that is OK
- Why do we need AI? - Information and data are increasing, we need help distilling
- When implementing and using AI educate users of the risks and look how it was created (i.e. AI Model Card)
- There are many examples of how AI can help clinical and administrative teams
  - Look at problem space and ROI
  - May not save time or make extra money but may decrease burden of the provider/care team and that is also OK

QUESTIONS & ANSWERS