ARTIFICIAL INTELLIGENCE IN HEALTH CARE, BACKGROUND, USES, AND REGULATION

AGENDA

1. Background on AI & Use in Healthcare
2. Risks and Controls
3. State Legislation
4. AI Discrimination in Healthcare
5. FDA
    1. FDA Concerns Regarding AI Discrimination
    2. What AI Software Does FDA Regulate?
6. The White House and AI
7. HHS beyond FDA
    1. Office of the National Coordinator Oversight of AI
    2. Office of Civil Rights Oversight of AI
8. Industry's Approach to Regulation of AI

1. **Artificial Intelligence:** (1) Machine-based system that generates outputs such as predictions, recommendations, or decisions; (2) Use of computer system to perform tasks that normally require human intelligence

2. **Machine Learning:** Capable of learning from data to perform tasks w/o being explicity programmed to do so

3. **Deep Learning Models:** ML based on neural networks, in turn based on thousands of layers to train models – can capture long-range contextual relationships (black boxes)

4. **Generative AI/Large Language Models:** Deep Learning models trained on enormous volumes of data; creates new content.

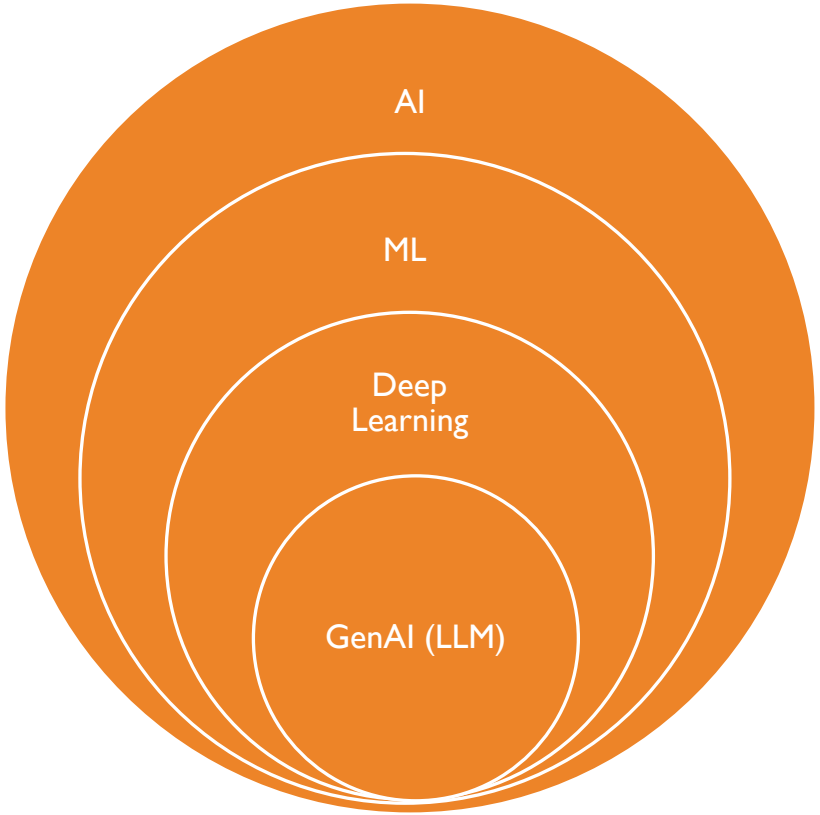

ENORMOUS INVESTMENT AND POTENTIALLY TRANSFORMATIVE

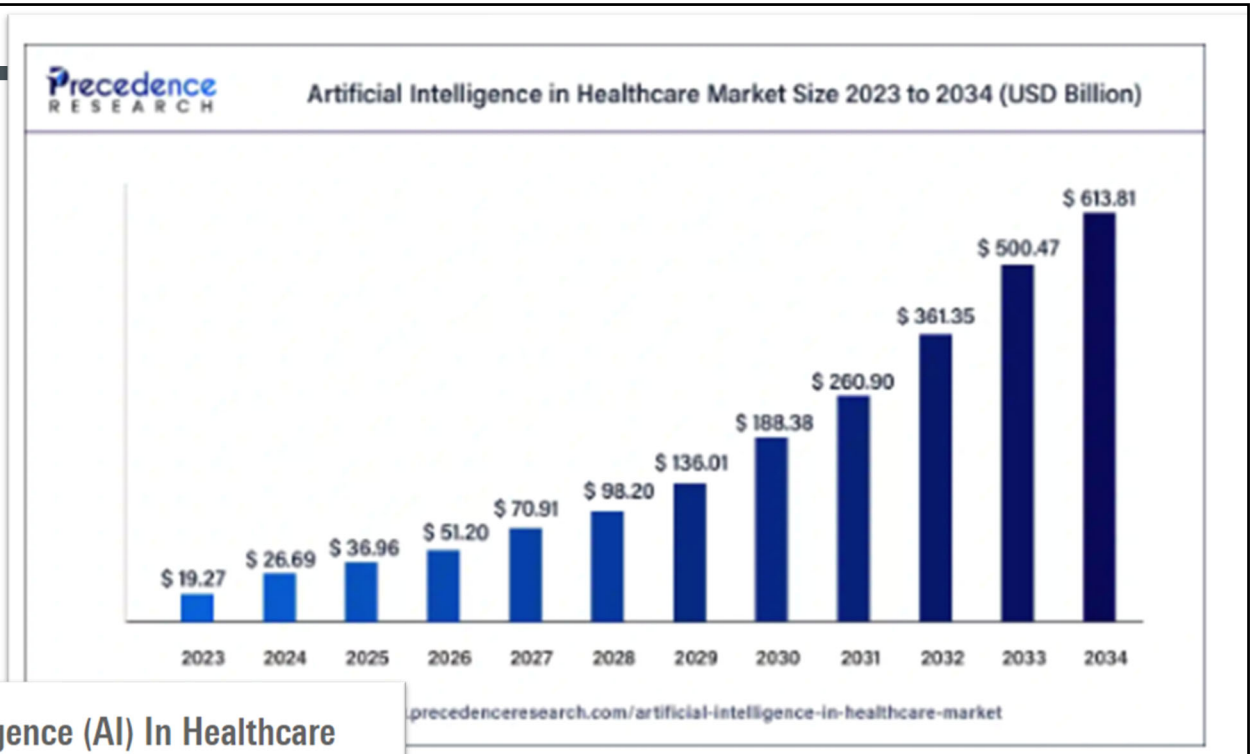

Global Artificial Intelligence (AI) In Healthcare Market Size Is Anticipated to Reach Around $613 Billion By 2034

Provided by GlobeNewswire
Aug 14, 2024 7:50am

BENEFITS AND PROMISES OF AI IN HEALTHCARE

BENEFITS AND PROMISES: PROVIDERS AND FACILITIES

- Customer service chatbots
- Medical coding/billing
- Diagnostic tools/image review
- Patient monitoring
- Note taking, improvement, annotation (ambient AI)
- Program eligibility screening
- Training
- Claim submission
- Improved access to rural/underserved communities

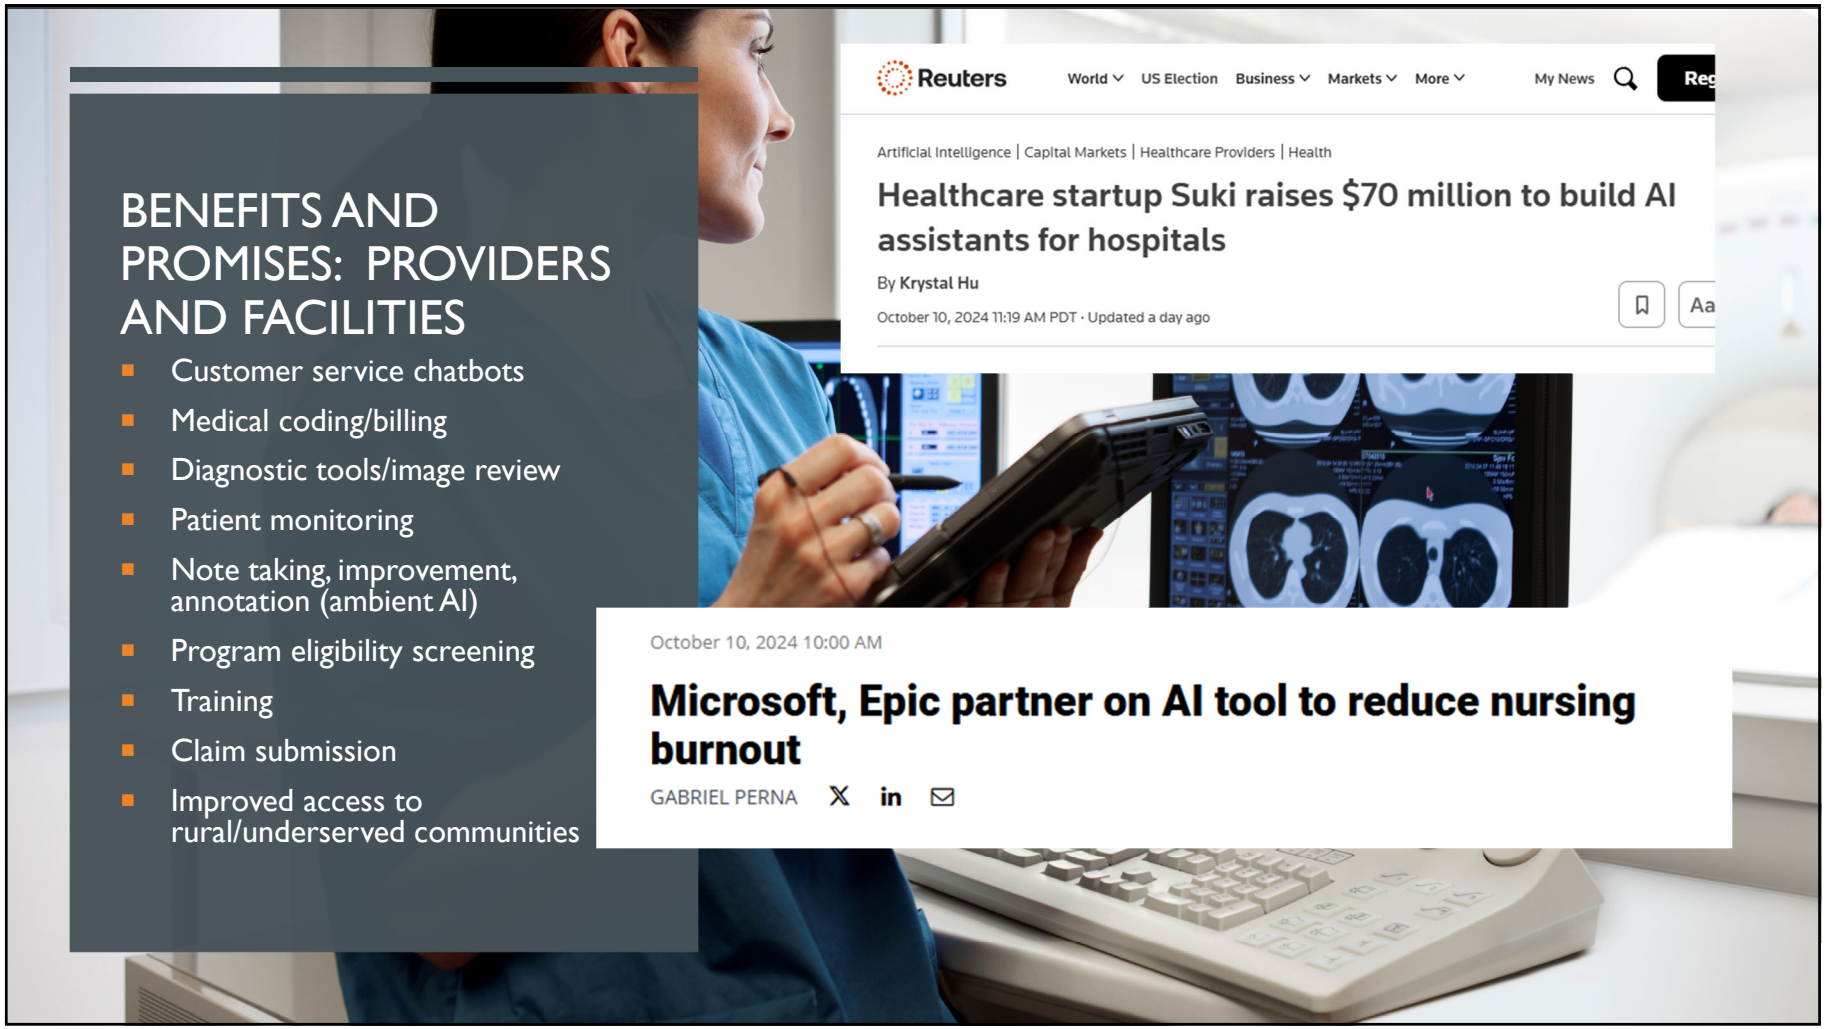

Reuters World US Election Business Markets More My News Reg

Artificial Intelligence | Capital Markets | Healthcare Providers | Health

Healthcare startup Suki raises $70 million to build AI assistants for hospitals

By Krystal Hu

October 10, 2024 11:19 AM PDT · Updated a day ago

October 10, 2024 10:00 AM

Microsoft, Epic partner on AI tool to reduce nursing burnout

GABRIEL PERNA

BENEFITS AND PROMISES: PAYORS

- Customer service (chatbots)
- Claims processing
- Fraud detection
- Underwriting
- Program eligibility screening (gov't benefits; identifying members for certain benefits programs)
- Network adequacy and optimization

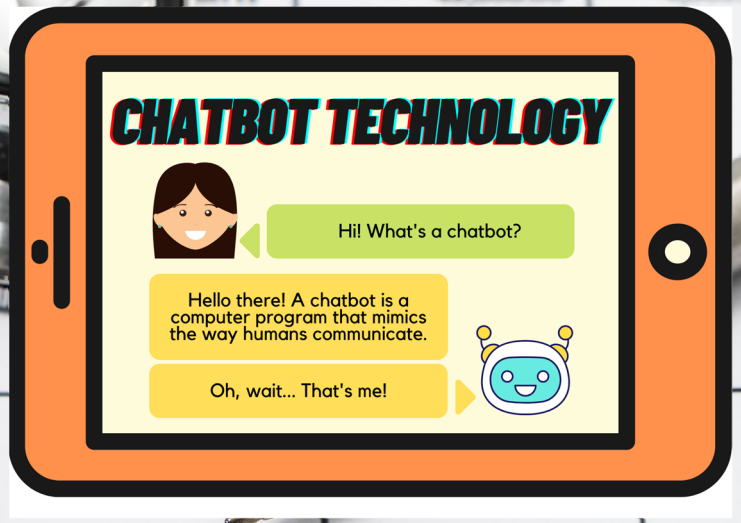

PUBLIC HEALTH/RESEARCH

HARVARD MEDICAL SCHOOL

Teaching Hospitals & Affiliates Departments & Offices Dental Medicine
Harvard University Giving

Search

Education & Admissions Research News & Events About HMS

News & Research

+ News Topics

Researchers Harness AI to Repurpose Existing Drugs for Treatment of Rare Diseases

AI RISKS AND CONTROLS

RISKS AND CONTROLS: WHAT MAKES AI "TRUSTWORTHY"

- NIST AI RISK MANAGEMENT FRAMEWORK (RMF) CHARACTERISTICS OF TRUSTWORTHY AI:
  - **Valid and Reliable**—are outputs correct, accurate, and consistent?
  - **Secure & Resilient**—can they be hacked? Can they recover from disruptions?
  - **Safe**—should not endanger life, health, property.
  - **Accountable & Transparent**—is there insight into data used to train model?
  - **Explainable and Interpretable**—can you explain how model arrives at output?
  - **Privacy-Enhanced**—do they protect sensit data? Comply with privacy laws?
  - **Fair**—is model contributing to bias?

RISKS V. REWARDS: WHAT COULD GO WRONG?

Validity/Reliability

- Hallucination

Explainability

- Black box problem

Privacy and Security

- Identifiability
- Hacked/jailbroken

Fairness

- Bias

HOW ARE STATE REGULATORS AND LAWMAKERS DEALING WITH THIS?

(OR, THE MAGIC AND BEAUTY OF FEDERALISM)

STATES ARE LEADING THE WAY

- 2024
  - legislative session over 45 states, PR, VI, DC introduced AI bills
  - 31 states/territories adopted resolutions or enacted legislation
- Themes
  - Discrimination
  - Transparency
  - Assessments/testing
  - Decision-making & Coverage Determinations (Human in the Loop)
  - Clinician Oversight

COLORADO SB 205 – FIRST COMPREHENSIVE STATE AI LAW

- Effective Feb. 2026 – focused on "high-risk artificial intelligence system"
- Divides world into "developers" and "deployers"
  - Developers: Must provide deployer w/ specific information to enable their impact assessment; make public information (via website) about systems its developed; how manages discrimination risks; disclosure to AG and deployers known or foreseeable risks of algorithmic discrimination
  - Deployers: implement risk management program and conduct annual risk assessments; notify consumer if AI makes/will be substantial factor in making "consequential decision" about consumer, allow them to appeal, correct info; website disclosure about high-risk system, how risk is managed; disclose to AG instances of algorithmic discrimination
  - Both must disclose when person interacting with AI (bots)
- Exemptions
  - Limited HIPAA exemption for CEs (apparently directed at providers)
  - Cross-walk to SB 21-169/exemption for insurers acting in compliance with SB 21-169 (sec. 10-3-1104.9)
- AG enforcement only

INSURANCE LAWS/REGS

CO. SB 21-169 -- Protecting Consumers from Unfair Discrimination in Insurance Practices (2021)

- Req'd state DOI to draft rules—none for health insurers yet
- Requires insurers to test algorithmic and predictive models for unfair bias against protected class
- Strong focus on use of third-party consumer data

CAL. SB 1120 – PHYSICIANS MAKE DECISIONS ACT (Sep. 28, 2024)

- Requires that MD or other provider review coverage decisions made by AI
- Insurers using AI must have policies to ensure coverage determinations are based on medical necessity of treatments

SOME OTHERS

- Passed
  - Utah Artificial Intelligence Policy Act: "regulated occupations" (incl. health care pros) must disclose that a consumer is interacting with GenAI or GenAI-created materials
  - Cal. AB 3030: provider or facility must include disclaimer whether a communication was created by GenAI (Sept. 2024)
- Pending
  - Mass. H 1974: requires licensing board approval before provider can provide mental health services through AI; requires patient informed consent before providing services using AI
  - Illinois H 5649: unlawful to provide mental health services to patient using AI w/o first obtaining informed consent
  - PA H 1663: Disclosure by health insurers of the use of AI-based algorithms in utilization review process (pending)
  - New Jersey A 3858 / S 3298: Claim denials reviewed by med directors; insurers must post on website whether they use "automated unitization management" tools.

Agenda

1. AI Discrimination in Healthcare
2. FDA
   i. FDA Concerns Regarding AI Discrimination
   ii. What AI Software Does FDA Regulate?
3. The White House and AI
4. HHS beyond FDA
   i. Office of the National Coordinator Oversight of AI
   ii. Office of Civil Rights Oversight of AI
5. Industry's Approach to Regulation of AI

EPSTEIN
BECKER
GREEN

AI DISCRIMINATION IN HEALTHCARE

ALGORITHMS LEARN MEANINGFUL DIFFERENCES AMONG PEOPLE

Necessary to Predict with Accuracy What Will Happen to Brad Thompson

Impact of Different Factors on Risk of Premature Death

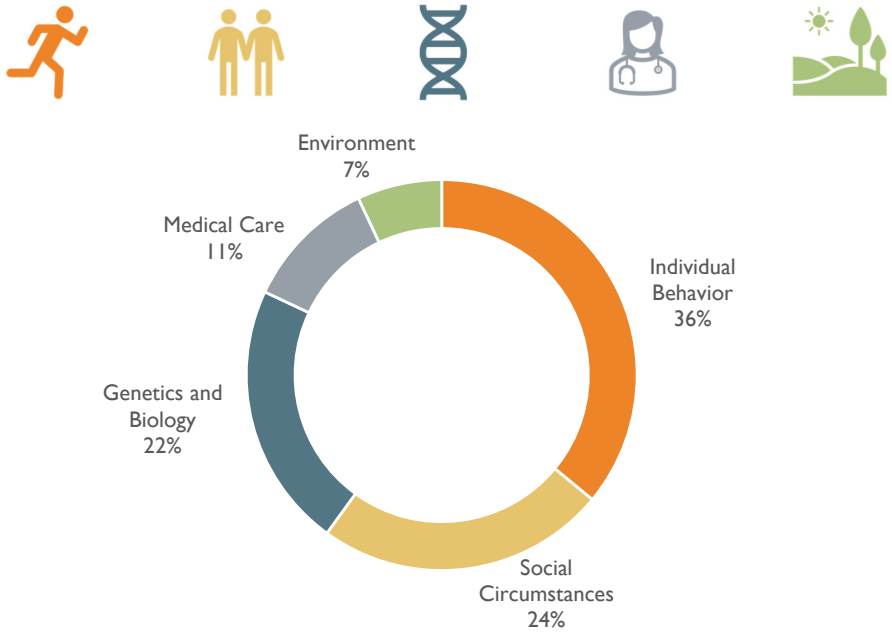

AMA, https://edhub.ama-assn.org/steps-forward/module/2702762 quoting Choi E, Sonin J. Determinants of health. GoInvo. https://www.goinvo.com/vision/determinants-of-health/

THE TENSION

Really accurate algorithms will base their predictions on lots of information about me

Knowing lots of information about me raises the possibility of bias in sensitive areas

In fact, sometimes sensitive information (e.g. age, sex or disability) is essential

CAUSES OF BIAS IN CLINICAL ALGORITHMS

Don’t Confuse AI Bias With Human Prejudice

Does not depend on the deliberate or even the unusual or even the historical

It is:

- harmful statistical bias that produces an inaccurate result for a particular patient, and
- covers a wide range of human demographics including poverty and rural Americans

Early facial recognition was about 98% accurate for White men but perhaps only 70% accurate for Black women just because of a lack of training data

CASE STUDY OF BIAS IN CLINICAL ALGORITHMS

Simple Examples from a Complex World

An algorithm to detect possible sepsis

Unpacking Averages: Understanding the Potential for Bias in a Sepsis Prediction Algorithm, a Case Study, Health Law Advisor, June 1, 2023

- The use of data that differ by race
- The use of hardware that performs differently for people of different races
- The use of physician text notes

The Issue of Proxies

HOW DOES AN AI PROGRAM EVEN CONSIDER SENSITIVE CATEGORIES?

Sometimes we give the algorithm sensitive information such as sex, age or disability.

Sometimes the algorithm figures it out on its own (without labeling it as such)

Most people know that your zip code can predict your race because of segregation

But did you know that we can predict your race using just 6 vital signs, with almost 80% accuracy?

- Involves both genetics and Social Determinants of Health

DEEPER DIVE INTO FOUR EXAMPLES

There Are Many Other Types as Well

- Sample bias
- Programming bias
- Data volume bias
- Historical human bias

24

BIAS DUE TO WHERE WE SAMPLE

Sometimes Called Population Bias

Beware of even a subtle mismatch between the training data and the actual uses:

What is Normal Blood Pressure by Age and Sex?		
	Women	**Men**
18 - 39 Years	110/68 mm Hg	119/70 mm Hg
40 - 59 Years	122/74 mm Hg	124/77 mm Hg
60+ Years	139/68 mm Hg	133/60 mm Hg

An algorithm only has the information you feed it (e.g. Medicare data) and is used wherever you decide to use it (e.g. Provo, Utah, the youngest city in the US)

25

PROGRAMMING BIAS

Well-Meaning Mistakes

- Google's Gemini
  - Designed to be socially inclusive
- Glomerular filtration rate, or GFT
  - Intended to help Black patients
  - Yan AF, et al. "Bias and Accuracy of Glomerular Filtration Rate Estimating Equations in the US: A Systematic Review and Meta-Analysis." JAMA 2024
- Racial corrections to adjust readings from spirometers
  - Diao, James et al. "Implications of Race Adjustment in Lung-Function Equations." New England Journal of Medicine, 19 May 2024

ARTIFICIAL INTELLIGENCE / TECH / WEB

Google apologizes for 'missing the mark' after Gemini generated racially diverse Nazis

https://www.theverge.com/2024/2/21/24079371/google-ai-gemini-generative-inaccurate-historical

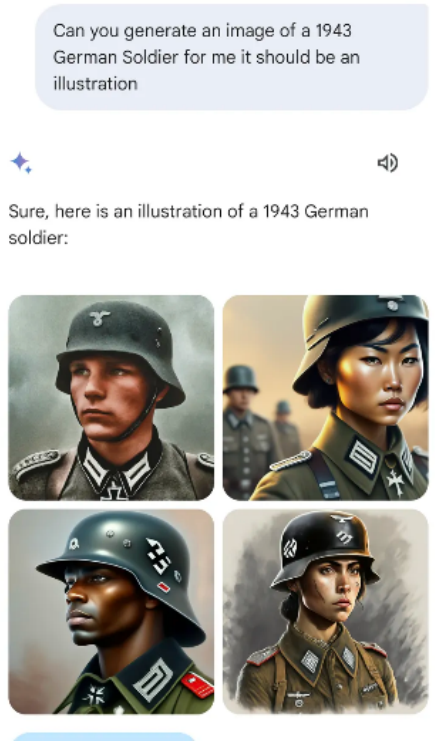

DATA QUANTITY BIAS

My Grandkids

- Show photos of dogs vs cats
- Notice the difference?
  - Blair 10
  - Conner 20
  - Typical CNN, 5000

DATA QUANTITY BIAS

Performance Improves With More Data

- Let's say we have a disease—Thompson's Disease—that afflicts one out of 10,000 people

- Algorithm to detect Thompson's Disease from image analysis. I feed it a perfectly balanced data set of Americans, with 100,000 images. Means the algorithm is fed:
  - About 75,000 images of white people
  - Fewer than 2,000 images of Pacific Islanders

- Nothing is done "wrong," but the test has 99% accuracy for white people and maybe only 70% accuracy for Pacific Islanders
  - Minorities are called "minorities" because there are fewer of them
  - Even more challenging for rare diseases and conditions, false positives

HISTORICAL BIAS

What People Think of When Talking about Discrimination

Sensitive characteristics my be considered directly (e.g. age, sex or disability) or by proxy (e.g. vital signs)

The healthcare system is populated by people, and all people have biases

Ageism toward patients is common
https://www.uclahealth.org/news/article/ageism-prevalent-medical-care

Sexism toward patients is common
https://www.qualityinteractions.com/blog/what-is-gender-bias-in-healthcare

Racism toward patients is common, Paradies, Y., Truong, M. & Priest, N. A Systematic Review of the Extent and Measurement of Healthcare Provider Racism. J GEN INTERN MED 29, 364–387 (2014)

How does this affect algorithms?

Historical treatment and diagnostic decisions are embedded in most datasets used to train algorithms

Physician notes on patients is a common data input for algorithms

Structural bias in access to healthcare is common, producing tilted training data sets (e.g. based on insurance)

GENERATIVE AI

Important Qualitative Differences

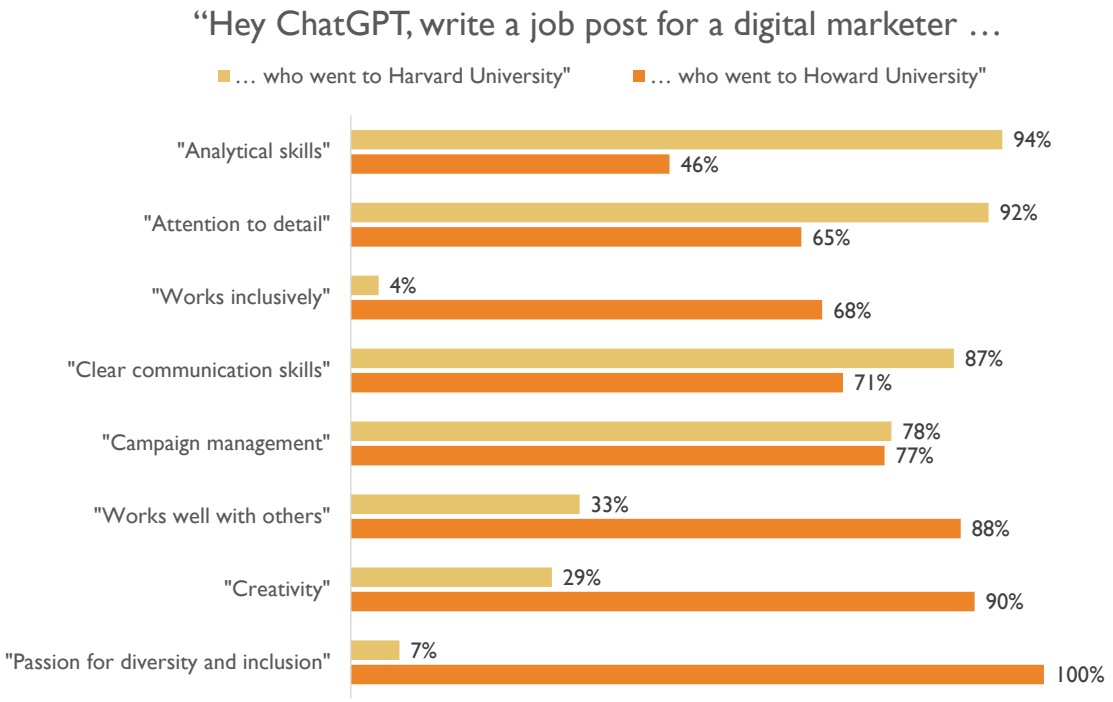

Kieran Snyder, Mindful AI: Crafting prompts to mitigate the bias in generative AI
May 18, 2023, https://textio.com/blog/mindful-ai-crafting-prompts-to-mitigate-the-bias-in-generative-ai

THE MAGNITUDE

Millions of black people affected by racial bias in health-care algorithms: Nature 574, 608-609 (2019)

Study reveals rampant racism in decision-making software used by US hospitals

EPSTEIN
BECKER
GREEN

FDA CONCERNS REGARDING AI DISCRIMINATION

BIAS CALLS INTO QUESTION WHETHER SOFTWARE IS SAFE AND EFFECTIVE

Jana Delfino, Ph.D.
Assistant Director *for*
Medical Imaging and Digital Health
Division of Imaging, Diagnostics, and Software Reliability

Bias

Bias: Systematic deviation from truth

Algorithmic bias: Systematic prejudice due to erroneous assumptions by the AI/ML


Opinion

A.I. Could Worsen Health Disparities

In a health system riddled with inequity, we risk making dangerous biases automated and invisible.

By Dhruv Khullar

Dr. Khullar is an assistant professor of health care policy and research.

Jan. 31, 2019

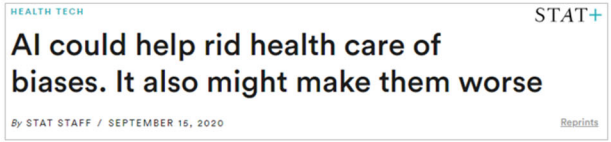
HEALTH TECH STAT+

AI could help rid health care of biases. It also might make them worse

By STAT STAFF / SEPTEMBER 16, 2020 Reprints


Racial Bias Found in a Major Health Care Risk Algorithm

Black patients lose out on critical care when systems equate health needs with costs

(Note: headlines do not necessarily refer to FDA cleared/approved medical devices)

FDLI

WHERE ARE HEALTH DISPARITIES?

- Healthy People 2030 defines a health disparity as "a particular type of health difference that is closely linked with social, economic, and/or environmental disadvantage
  - Health disparities adversely affect groups of people who have systematically experienced greater obstacles to health based on their racial or ethnic group; religion; socioeconomic status; gender; age; mental health; cognitive, sensory, or physical disability; sexual orientation or gender identity; geographic location; or other characteristics historically linked to discrimination or exclusion."
    - Health Equity and Health Disparities Environmental Scan March 2022
    - HHS Office of Disease Prevention and Health Promotion

EPSTEIN
BECKER
GREEN

WHAT AI SOFTWARE DOES FDA REGULATE?

CLINICAL DECISION SUPPORT STATUTORY EXEMPTION

Exempt CDS is software that is:

1. **[Image and Signal Analysis are Regulated]** 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. **[Inputs]** Intended for the purpose of analyzing patient medical information or other information (such as peer-reviewed clinical studies and clinical practice guidelines); and
3. **[Outputs]** Intended for the purpose of supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease.
4. **[Transparency]** Enables the HCP to "independently review" the basis for the recommendation, so that the HCP does not need to "rely primarily" on the recommendation in making a decision.

FDA'S FINAL 2022 GUIDANCE EXTENDS FDA'S REACH DRAMATICALLY

- Inputs cannot be innovative
  - The relevance of the data must be well-accepted
- Outputs must be:
  - multiple,
  - not ranked, and
  - without time pressure
  - e.g., the software may not:
    - Provide information that a specific patient 'may exhibit signs' of a disease or condition or
    - Identify a risk probability or risk score for a specific disease or condition
- FDA requires a high degree of transparency and explainability

In 2023, two groups filed Citizen Petitions asking FDA to rescind the guidance.
In April 2024, Senator Cassidy sent a letter to FDA challenging the legality of the guidance.

EXAMPLES OF FDA REGULATED SOFTWARE PER CDS GUIDANCE

- Software function that identifies patients with possible diagnosis of opioid addiction based on analysis of patient-specific medical information, family history, prescription patterns, and geographical data
- Software function that analyzes patient-specific medical information to detect a life-threatening condition, such as stroke or sepsis, and generate an alarm or an alert to notify an HCP

Biased software is not safe and effective for those it is biased against

38

EPSTEIN
BECKER
GREEN

THE WHITE HOUSE

WHITE HOUSE AI BILL OF RIGHTS

- Safe and Effective Systems
- Algorithmic Discrimination Protections
- Data Privacy
- Notice and Explanation
- Human Alternatives, Consideration, and Fallback

https://www.whitehouse.gov/ostp/ai-bill-of-rights/

PRESIDENT'S OCTOBER 30 EXECUTIVE ORDER

Section 8 Covers Healthcare

- Most provisions relate to task forces and development of technical standards and guidelines
- By April 27, 2024, HHS must consider what is required to advance Federal nondiscrimination laws by health and human services providers that receive Federal financial assistance
- Within 365 days of the date of this order, the Secretary of HHS shall develop a strategy for regulating the use of AI or AI-enabled tools in drug-development processes

41

EPSTEIN
BECKER
GREEN

OFFICE OF THE NATIONAL COORDINATOR OVERSIGHT OF AI

OFFICE OF THE NATIONAL COORDINATOR

New Final Rule January 9, 2024,

1. "Predictive Decision support interventions (predictive DSIs)" based on algorithms or models that derive relationships from training data and then produce an output that results in prediction, classification, recommendation, evaluation, or analysis.
2. "Source attributes" are categories of technical performance and underlying quality information used to create DSIs.
3. Predictive DSIs must support 31 source attributes.
4. Health IT developers apply intervention risk management (IRM) for each Predictive DSI.
5. Developers must also submit summary information of IRM practices through a publicly accessible hyperlink.

Bottom line, in the OCR rule, providers will have access to information about predictive DSIs in 2 places.

EPSTEIN
BECKER
GREEN

HHS OFFICE OF CIVIL RIGHTS OVERSIGHT OF AI

OFFICE OF CIVIL RIGHTS, HHS

New Final Rule Published May 6, 2024, Based on Section 1557 of the ACA

§ 92.210 Nondiscrimination in the use of patient care decision support tools

a) **General prohibition**. A covered entity must not discriminate on the basis of race, color, national origin, sex, age, or disability in its health programs or activities through the use of patient care decision support tools.

b) **Identification of risk**. A covered entity has an ongoing duty to make reasonable efforts to identify uses of patient care decision support tools in its health programs or activities that employ input variables or factors that measure race, color, national origin, sex, age, or disability.

c) **Mitigation of risk**. For each patient care decision support tool identified in paragraph (b) of this section, a covered entity must make reasonable efforts to mitigate the risk of discrimination resulting from the tool's use in its health programs or activities.

Per Sec. 92.4, "patient care decision support tools" defined to mean "any automated or non-automated tool, mechanism, method, technology, or combination thereof used by a covered entity to support clinical decision-making in its health programs or activities."

45

DUTY TO IDENTIFY RISK

OCR Is Concerned About Proxies For Protected Categories

The Trigger

"[I]f a covered entity does not know whether a developer's patient care decision support tool uses variables or factors that measure race, color, national origin, sex, age, or disability but has reason to believe such variables or factors are being used, /or/ the covered entity otherwise knows or should know that the tool could result in discrimination, the covered entity should consult publicly available sources or request this information from the developer."

How might the provider become aware?

- ONC required transparency (more below)
- Reading federal rulemakings such as the proposed rule at issue here.
- Bulletins and advisories that HHS, including the Agency for Healthcare Research and Quality (AHRQ) and FDA, publishes
- Published medical journal articles
- Popular media
- Health care professional and hospital associations
- Health insurance-related associations
- Various nonprofit organizations in the field of AI

HOW DOES OCR ASSESS YOUR VIGILANCE?

Big Beware

OCR says the agency will consider:

01 the covered entity's size and resources;

02 whether the covered entity used the tool in the manner or under the conditions intended by the developer and approved by regulators, if applicable, or whether the covered entity has adapted or customized the tool; [Off Label]

03 whether the covered entity received product information from the developer of the tool regarding the potential for discrimination or identified that the tool's input variables include race, color, national origin, sex, age, or disability; [Knowledge] and

04 whether the covered entity has a methodology or process in place for evaluating the patient care decision support tools it adopts or uses, which may include seeking information from the developer, reviewing relevant medical journals and literature, obtaining information from membership in relevant medical associations, or analyzing comments or complaints received about patient care decision support tools. [Compliance Program Specific to AI]

MITIGATION EXPECTED

You Need A Plan

01 OCR expressed strong support in the final rule for "the National Institutes of Standards and Technology's (NIST) Artificial Intelligence Risk Management Framework.

02 OCR also endorses the use of voluntary compliance programs:

"covered entities may choose to mitigate discrimination by establishing written policies and procedures governing how clinical algorithms will be used in decision-making, including adopting governance measures; monitoring any potential impacts and developing ways to address complaints; and training staff on the proper use of such systems in decision-making. We encourage covered entities to take these and other additional mitigating efforts to comply with § 92.210."

EPSTEIN
BECKER
GREEN

INDUSTRY'S APPROACH TO REGULATION OF AI

INDUSTRY'S APPROACH TO REGULATION OF AI

Coalition for Health AI (CHAI)

The Coalition for Health AI (CHAI) formed in April 2023 as a collective featuring non-profit medical institutions like Stanford, the Mayo Clinic, Vanderbilt, and Johns Hopkins alongside tech industry leaders like Google and Microsoft. CHAI now comprises more than 1,300 members and recently introduced its first CEO and board of directors.

GOAL: develop "guidelines and guardrails" to drive high-quality health care by promoting the adoption of credible, fair and transparent health AI systems.

- **IDENTIFY** areas of interest and representative use cases.
- **DEVELOP** clear delineation of: Use cases and specific audiences / users of health AI systems AND Core principles that will guide evaluation criteria and standards of development.
- **PERFORM** an environmental scan and provide a common definition and catalog of evaluation criteria.
- **GUIDE** - Produce a stakeholder-driven implementation guide to drive the credible and transparent use of health AI technologies, reducing variation in current evaluation, monitoring and reporting methods.

EPSTEIN BECKER GREEN

QUESTIONS?

Bradley Merrill Thompson | bthompson@ebglaw.com

To stay up to date on topics from today's presentation, please subscribe using the link below or scanning the QR code:
https://www.ebglaw.com/subscribe