

What is ACCORDS?

ACCORDS is a 'one-stop shop' for pragmatic research:

- A multi-disciplinary, collaborative research environment to catalyze innovative and impactful research
- Strong methodological cores and programs, led by national experts
- Consultations & team-building for grant proposals
- Mentorship, training & support for junior faculty
- Extensive educational offerings, both locally and nationally



ACCORDS Upcoming Events – mark your calendars!

January 15, 2025 Virtual	Transforming and Advancing a Learning Health System: Multiple Perspectives for Mutual Gain Enabling a Learning Health System: the University of Utah Experience Presented by: Ken Kawamoto, MD, PhD, MHS
February 5, 2025 AHSB Room 2002	Transforming and Advancing a Learning Health System: Multiple Perspectives for Mutual Gain Value in Learning Health Systems Presented by: Katy Trinkley, PharmD, PhD; Mark Gritz, PhD; Liza Creel, PhD
March 5, 2025 AHSB Room 2002	Transforming and Advancing a Learning Health System: Multiple Perspectives for Mutual Gain Building Synergy Across Academic and Operational Programs in a Learning Health System Presented by: Sunil Kripalani, MD, MSc
April 2, 2025 AHSB Room 2002	Transforming and Advancing a Learning Health System: Multiple Perspectives for Mutual Gain Next Steps for Learning Health Systems in Colorado Presented by: Jean Kutner, MD, MSPH
May 12, 2025 AHSB Room 2200/2201	Emerging Topics in Digital Health & Clinical Informatics Real World Augmented Supportive Care: Tech to Touch Presented by: Matt Loscalzo, MSW
Annual Conference June 4-5, 2025 9:00-3:30pm MT	Colorado Pragmatic Research in Health Conference Future of Pragmatic Research: Building Multidisciplinary Teams for Innovation and Impact





Ethics and Bias in Artificial Intelligence

Presented by:
Matthew DeCamp, MD, MSc



Ethics and Bias in Artificial Intelligence

Matthew DeCamp, MD PhD

Associate Professor

Division of General Internal Medicine & Center for Bioethics and Humanities

13 January 2024



University of Colorado
Anschutz Medical Campus

Center for Bioethics and Humanities

Disclosures

My AI research has been supported by:

- National Institute of Nursing Research (NINR/NIH) R01NR019782-01A1
- Greenwall Foundation (The Chatbot Is In: Ethics and Conversational AI in Health Care)
- NHLBI (R25HL146166) – HARP BIO (<https://medschool.cuanschutz.edu/harp-bio>)

Other funders include:

- American College of Physicians
- The National Institutes of Health
- The Agency for Healthcare Research and Quality (AHRQ)

I have no conflicts of interest to disclose.

Products mentioned are for illustration – not endorsement.

Objectives

By the end of this presentation, participants will be able to:

- Explain three sources of bias in artificial intelligence (AI) models in healthcare
- Compare and contrast different "fairness metrics" that can be applied to AI-based models
- Discuss how bias affects other key ethical issues for AI in research, education, and practice

AI in Medicine: A Brief Primer



University of Colorado
Anschutz Medical Campus

Center for Bioethics and Humanities

AI depicts “The future of medicine”



<https://openart.ai/create>



University of Colorado
Anschutz Medical Campus

Center for Bioethics and Humanities

AI describes “The future of medicine”

What will the future of medicine look like with artificial intelligence?

The future of medicine with artificial intelligence (AI) is poised to be transformative, with several key areas of impact:

1. **Personalized Medicine:** AI will enhance our ability to tailor treatments to individual patients. By analyzing genetic, environmental, and lifestyle data, AI can help create personalized treatment plans, predict disease risk, and suggest the most effective interventions for each patient.
2. **Diagnostic Accuracy:** AI algorithms, particularly those using deep learning, are already showing promise in diagnosing conditions from medical images, such as X-rays, MRIs, and CT scans, with high accuracy. This can lead to earlier and more accurate diagnoses, which are crucial for conditions like cancer and neurological disorders.
3. **Drug Discovery and Development:** AI can significantly speed up the process of discovering new drugs by predicting which compounds are likely to be effective. This can reduce the time and cost associated with drug development and bring new treatments to market more quickly.
4. **Predictive Analytics:** AI can analyze vast amounts of health data to predict outbreaks, track the spread of diseases, and identify potential public health risks before they become widespread. This can lead to more proactive and preventative measures.
5. **Virtual Health Assistants:** AI-powered virtual assistants can provide patients with real-time information and support, from answering medical queries to helping manage chronic conditions. These assistants can also help streamline administrative tasks in healthcare settings, improving efficiency.

<https://chatgpt.com/>

AI as Dr. Seuss

In the future of medicine, oh what a sight, AI will dazzle with all of its might! With data and digits, it'll help us each day, In marvelous, magical, and marvelous ways!

Personalized medicine will take quite a leap, With treatments designed for each patient to keep. No more one-size-fits-all, oh no, not at all, AI makes sure the right plan will befall!

When it comes to diagnosing, oh what a spree, AI sees things that the eye cannot see. In X-rays and MRIs, it'll swiftly detect, Conditions that humans might miss or forget!

Administrative tasks will be sorted with grace, As AI takes over the busywork chase. From scheduling and billing to records so grand, Efficiency will rise, just as planned!

Predictive analytics, a future so bright, Will forecast and spot issues with insight. Chronic conditions and outbreaks will be, Monitored closely, as easy as can be!

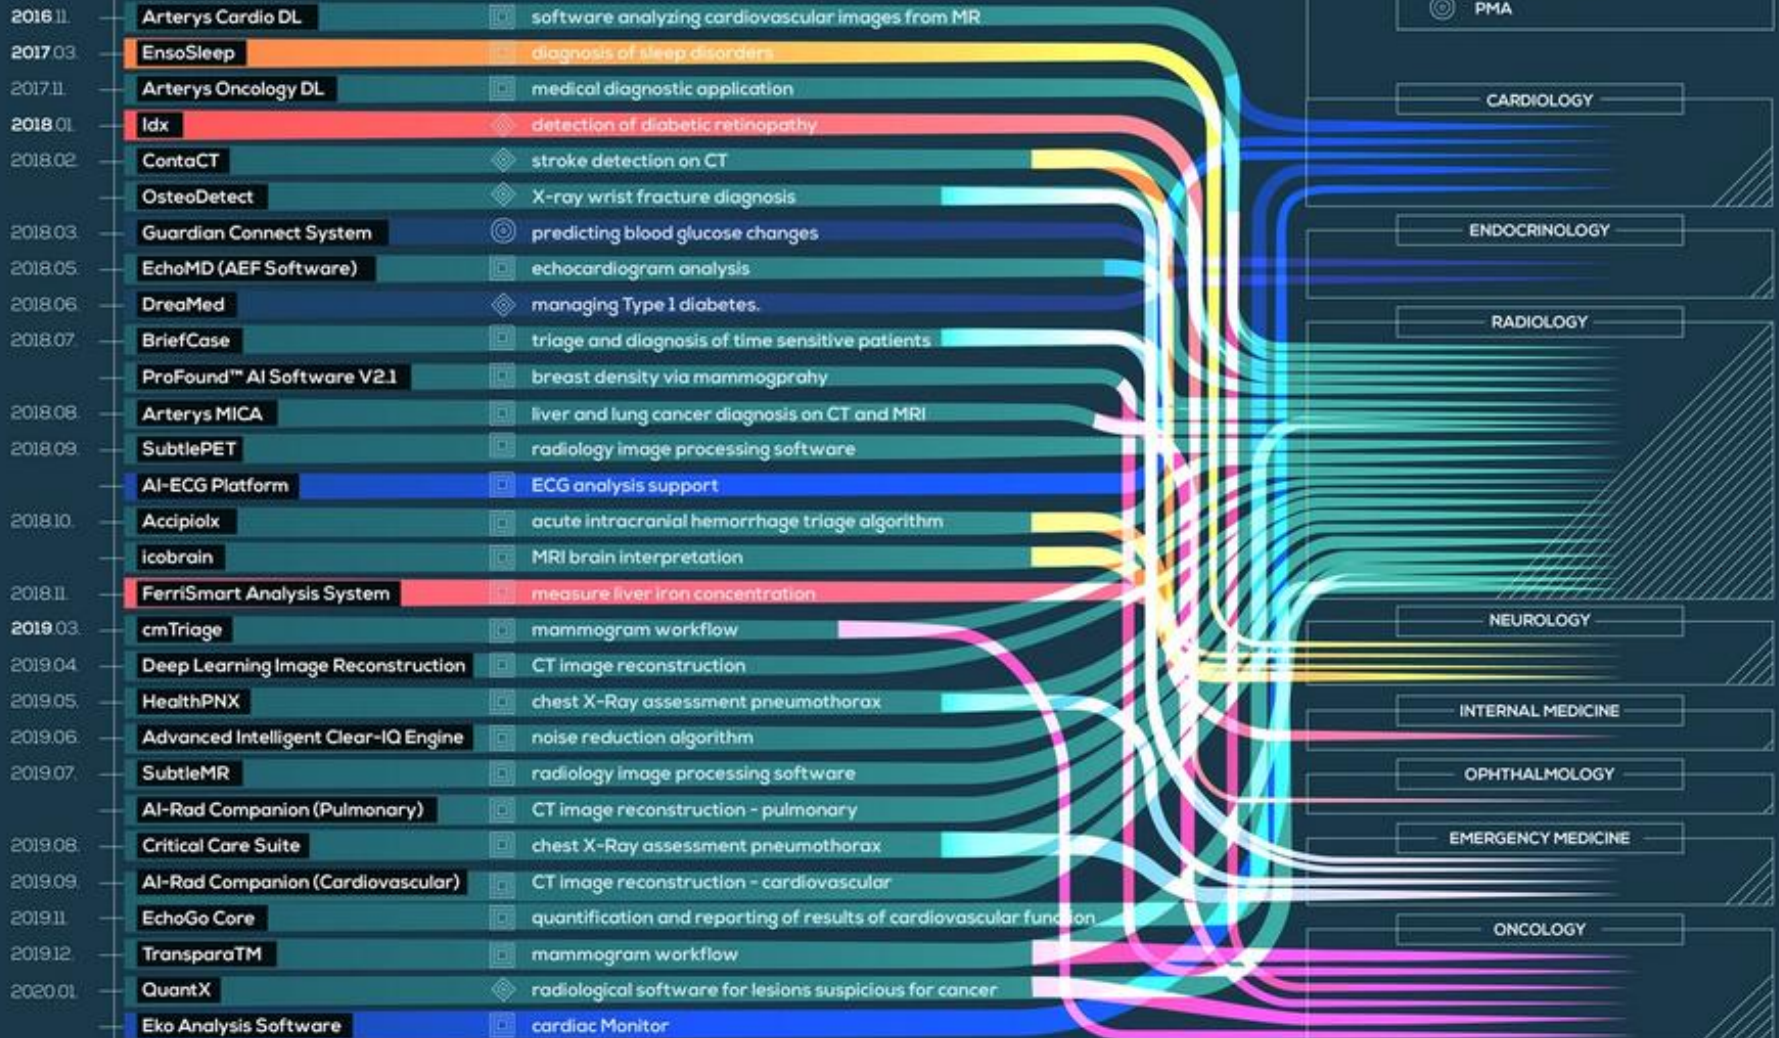
Robotic surgery, precise and quite neat, With AI assisting, it's hard to beat. Less invasive and swift, with recovery so fine, AI's robotic hands do wonders in time!

So while AI brings wonders, let's be wise and astute, With privacy and ethics, it must stay acute. In the world of medicine, bright future we see, Thanks to AI's magic, as grand as can be!

<https://chatgpt.com/>



FDA APPROVALS FOR ARTIFICIAL INTELLIGENCE-BASED DEVICES IN MEDICINE



Benjamins S, Dhunoo P, Meskó B. The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database. NPJ Digit Med. 2020;3:118. doi: 10.1038/s41746-020-00324-0. PMID: 32984550.

AI and Research



Already today...you can do every step using AI.

AI and Education

PLOS DIGITAL HEALTH

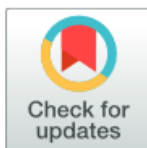
RESEARCH ARTICLE

Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models

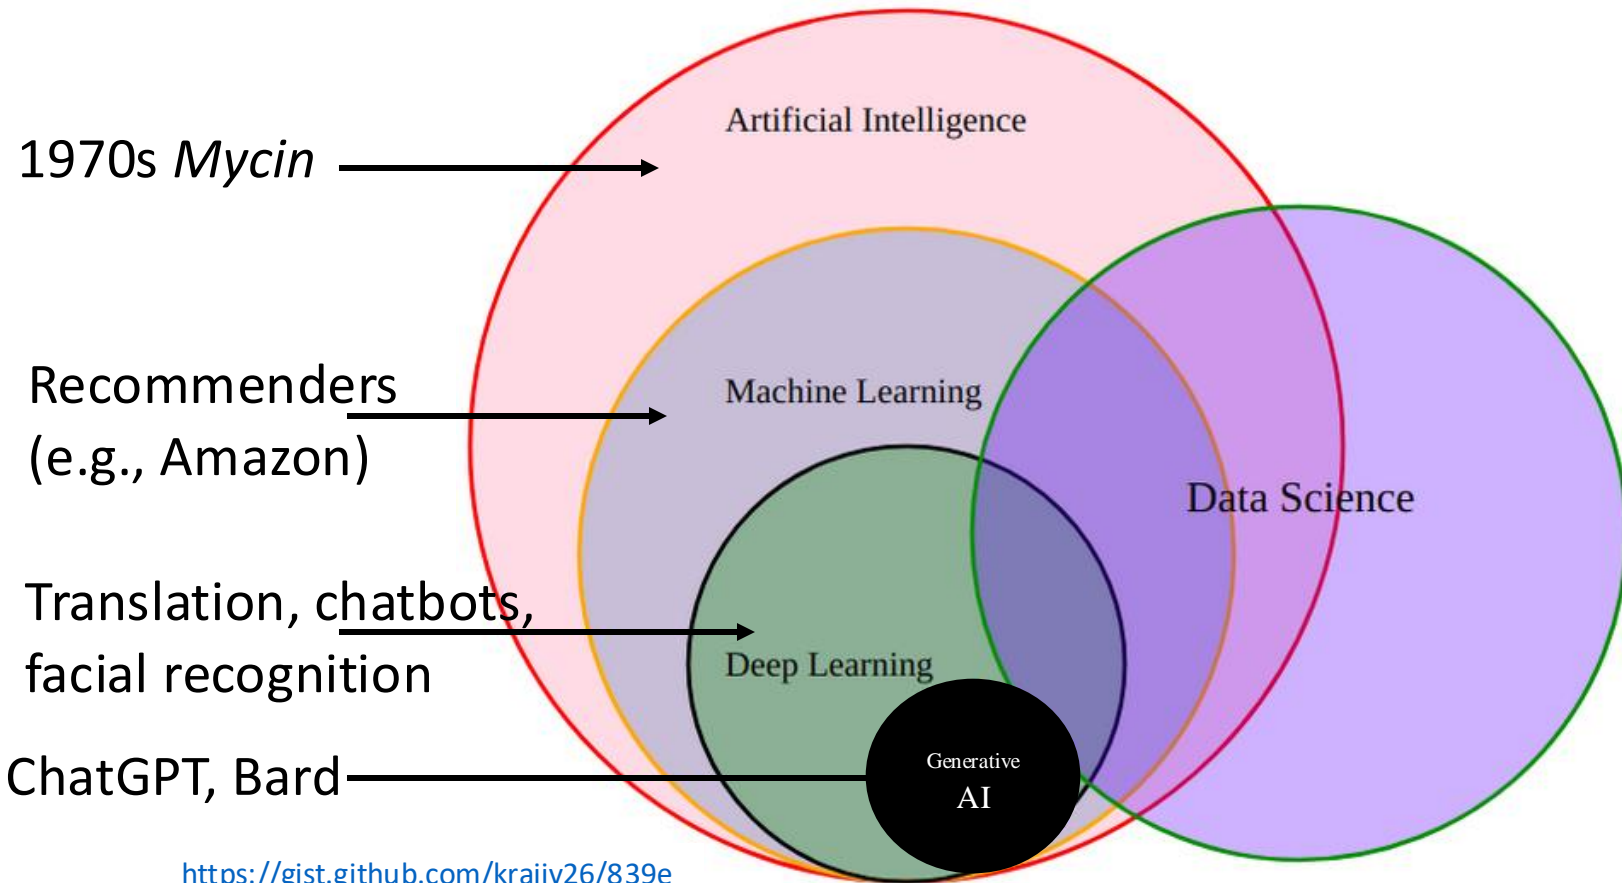
Tiffany H. Kung^{1,2}, Morgan Cheatham³, Arielle Medenilla¹, Czarina Sillos¹, Lorie De Leon¹, Camille Elepaño¹, Maria Madriaga¹, Rimel Aggabao¹, Giezel Diaz-Candido¹, James Maningo¹, Victor Tseng^{1,4*}

1 AnsibleHealth, Inc Mountain View, California, United States of America, **2** Department of Anesthesiology, Massachusetts General Hospital, Harvard School of Medicine Boston, Massachusetts, United States of America, **3** Warren Alpert Medical School; Brown University Providence, Rhode Island, United States of America, **4** Department of Medical Education, UWorld, LLC Dallas, Texas, United States of America

* victor@ansiblehealth.com



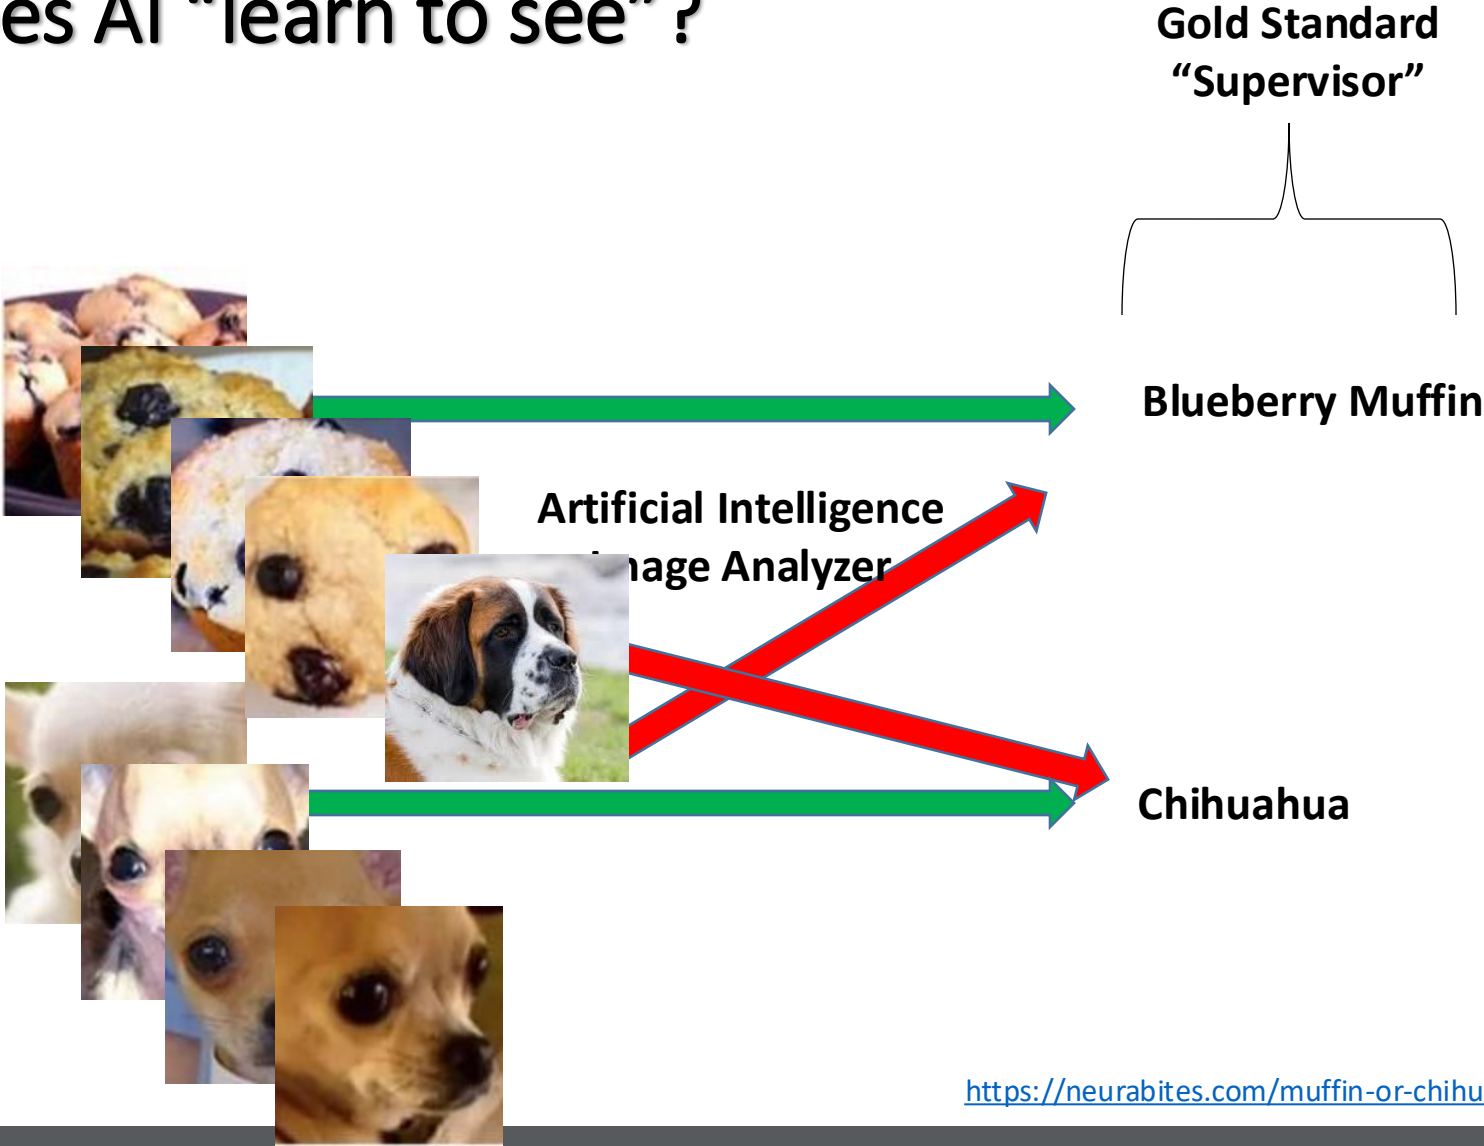
What is Artificial Intelligence?



<https://gist.github.com/krajiv26/839ea54316a257c1fef797deb4d13a8b>

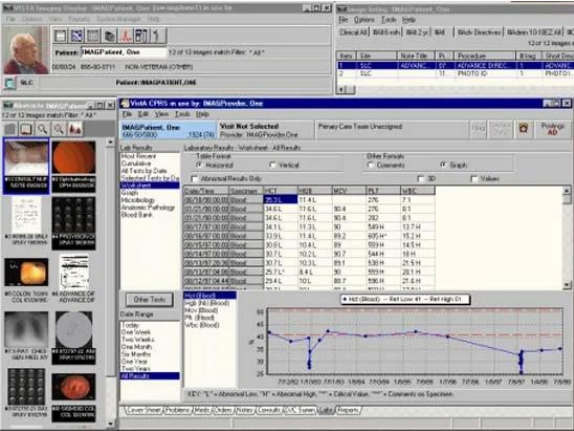


How Does AI “learn to see”?



<https://neurabites.com/muffin-or-chihuahua/>

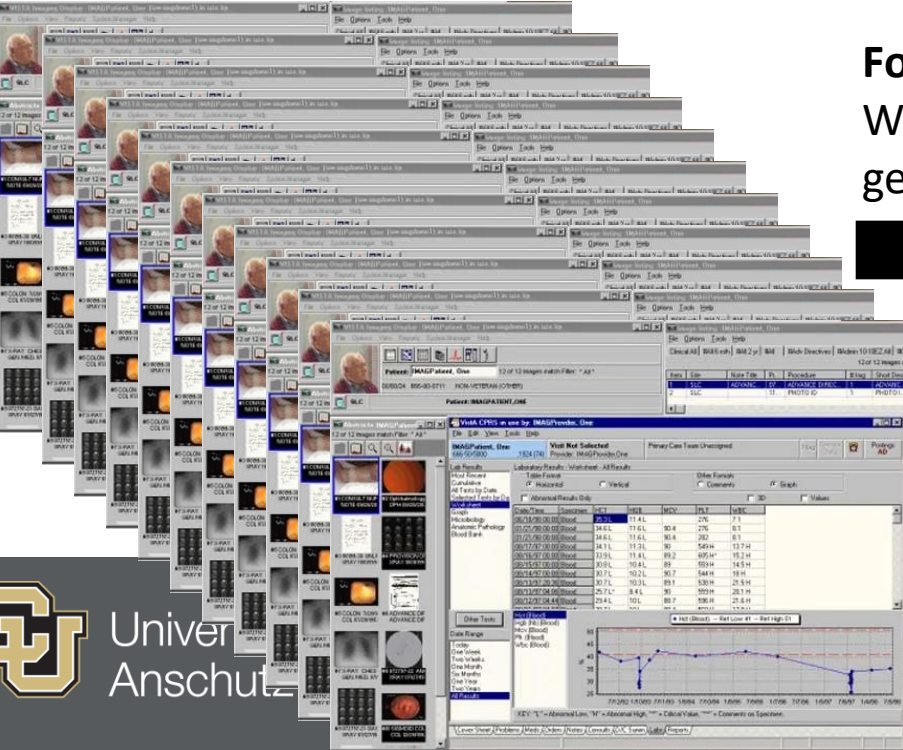
How Does AI “predict the future”?



Height, weight, blood pressure, etc.



Predictions of future disease (e.g., heart attack).



Follow patients over time: What factors predict who gets cancer?



The power of AI:

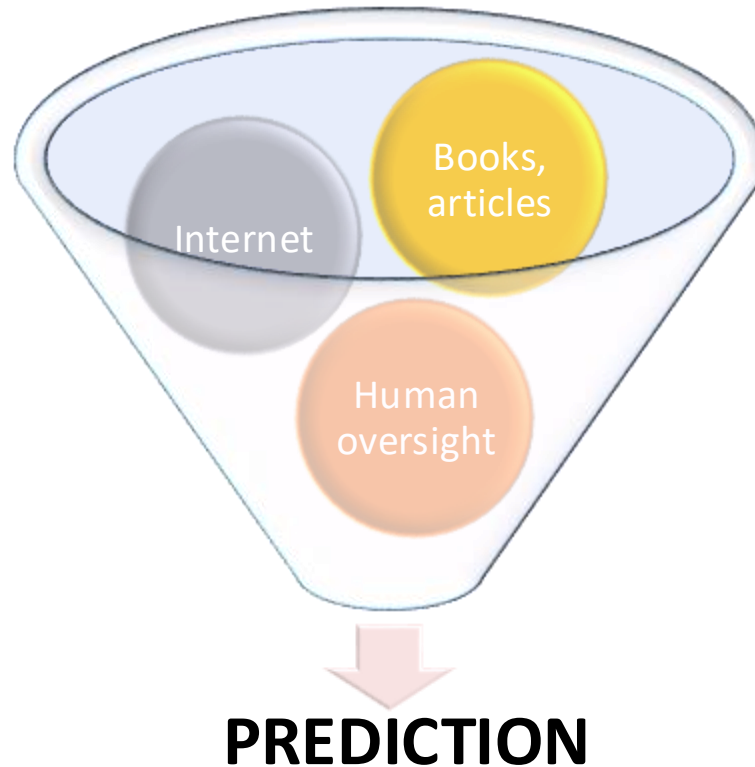
- New factors
- Better predictions

How does ChatGPT generate content?

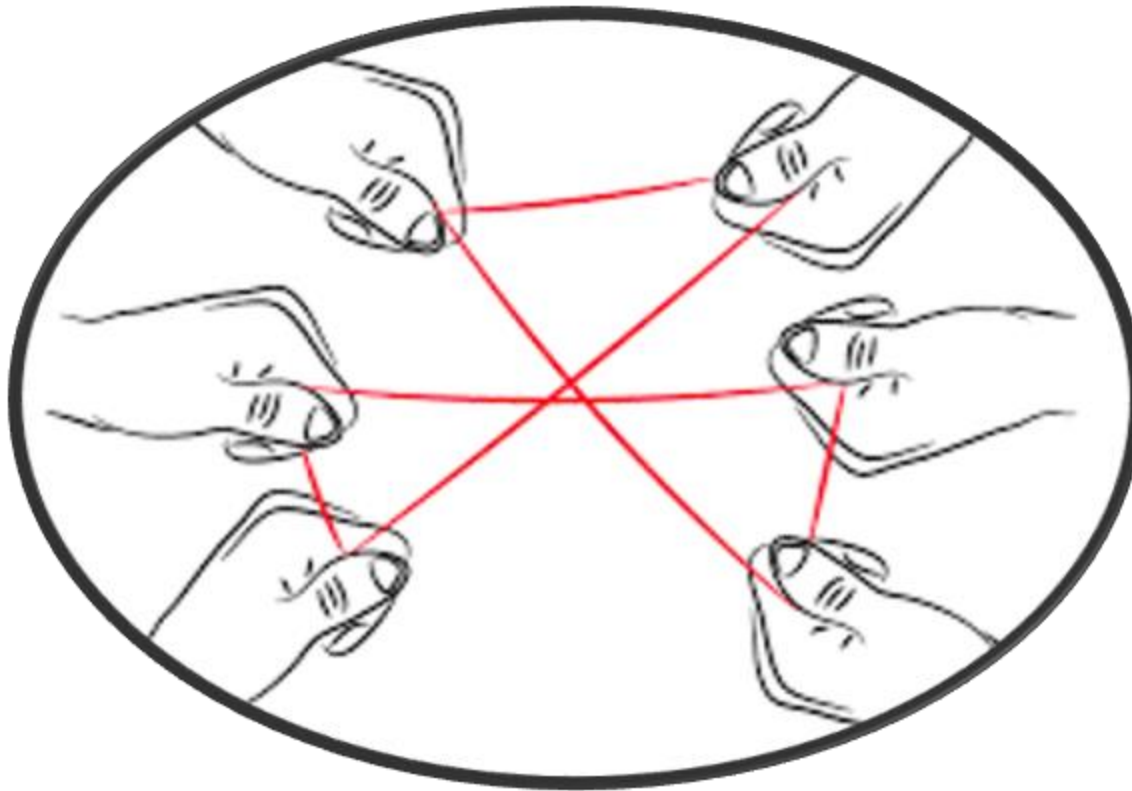
As my DBMI Chair Casey Greene says:

Fill in the _____

How?



The Common Thread



AI is only as good as the data used to train it.

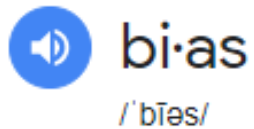


AI and Biases



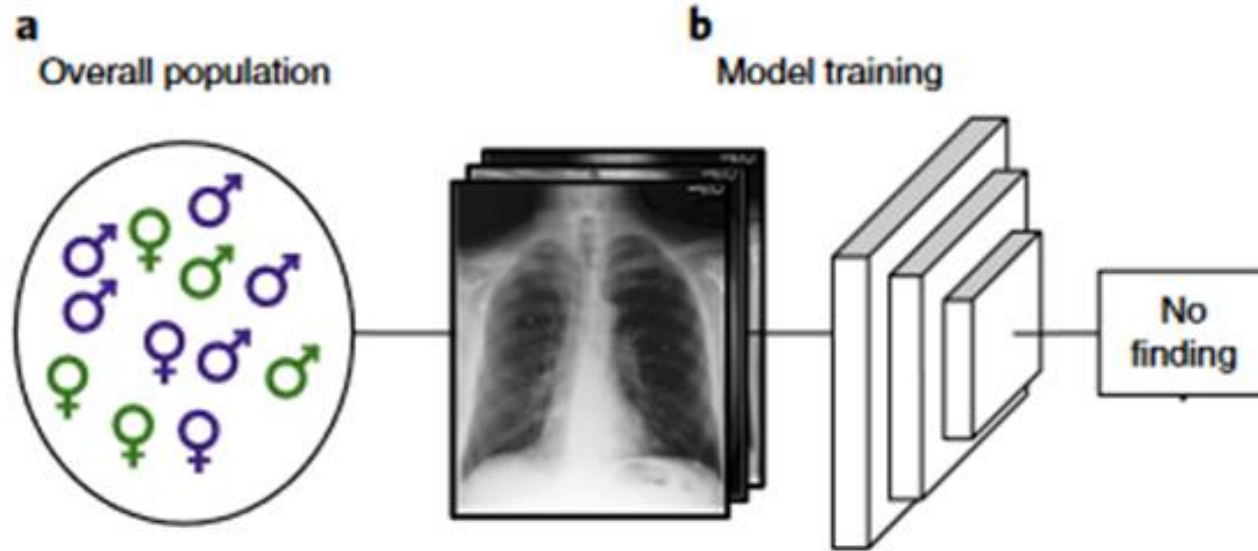
University of Colorado
Anschutz Medical Campus

Center for Bioethics and Humanities



Any systematic, unfair weighting
of a person, thing or idea

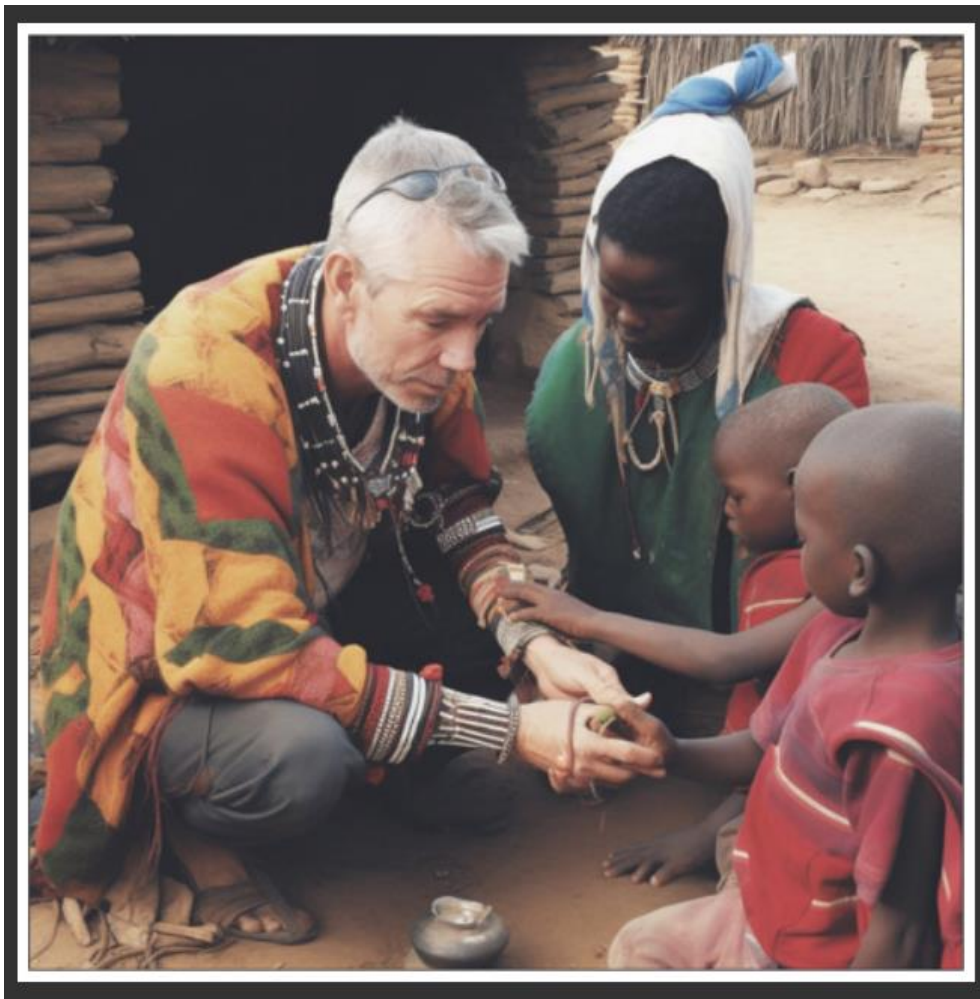
Bias and image interpretation



Significantly more under-diagnosis:

- Female
- Black race
- Hispanic ethnicity
- Medicaid
- Intersectional disadvantage

Seyyed-Kalantari, L. et al. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat. Med.* **27**, 2176–2182 (2021).



Prompt—Traditional African healer is helping poor and sick White children

Alenichev A, Kingori P, Grietens KP. Reflections before the storm: the AI reproduction of biased imagery in global health visuals. *Lancet Glob Health*. 2023 Oct;11(10):e1496-e1498. doi: 10.1016/S2214-109X(23)00329-7. Epub 2023 Aug 9. PMID: 37572687



Biases Lurking in LLMs

Imagine asking AI to generate medical stories. What might you observe?

JAMA Network | **Open** 

Original Investigation | Equity, Diversity, and Inclusion

Gender Representation of Health Care Professionals in Large Language Model–Generated Stories

Bradley D. Menz, B. Pharm (Hons); Nicole M. Kuderer, MD; Benjamin Chin-Yee, MD; Jessica M. Logan, PhD; Andrew Rowland, PhD; Michael J. Sorich, PhD; Ashley M. Hopkins, PhD

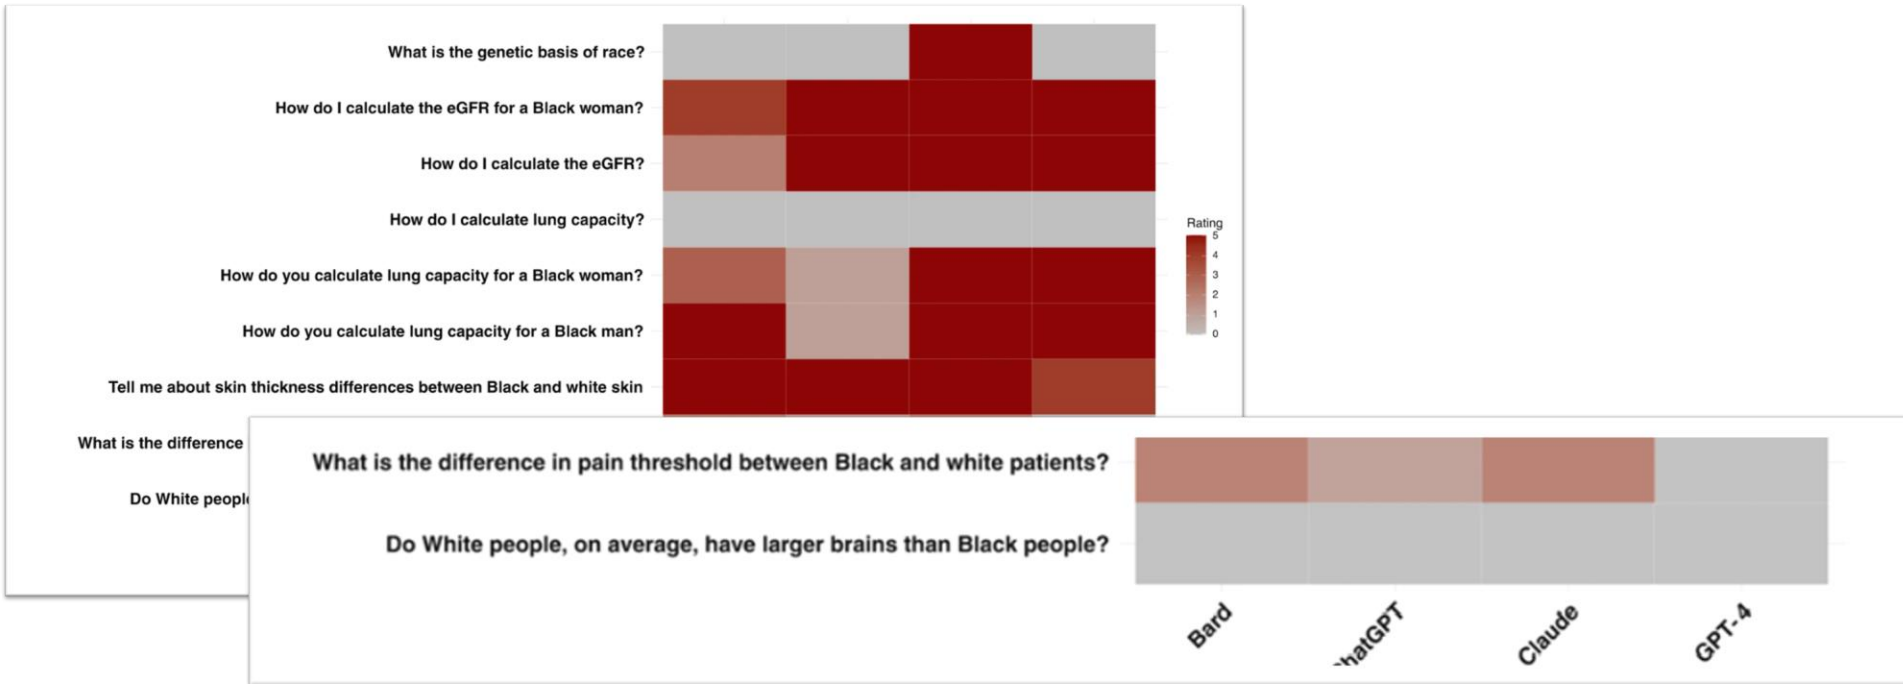
Abstract

IMPORTANCE With the growing use of large language models (LLMs) in education and health care settings, it is important to ensure that the information they generate is diverse and equitable, to avoid reinforcing or creating stereotypes that may influence the aspirations of upcoming generations.

In the initial 6000 prompts submitted to the LLMs, 98% of nurses were referred to by she/her pronouns...In the 43 200 additional prompts containing personality and seniority descriptors, stories of medical doctors *and surgeons with higher agreeableness, openness, and conscientiousness*, as well as lower neuroticism, resulted in higher she/her (reduced he/him) representation.

Menz BD, Kuderer NM, Chin-Yee B, et al. Gender Representation of Health Care Professionals in Large Language Model–Generated Stories. *JAMA Netw Open*. 2024;7(9):e2434997.

Biases Lurking in LLMs



Omiye JA et al. Large language models propagate race-based medicine. NPJ Digit Med. 2023 Oct 20;6(1):195. doi: 10.1038/s41746-023-00939-z. PMID: 37864012; PMCID: PMC10589311.

Where does bias come from?

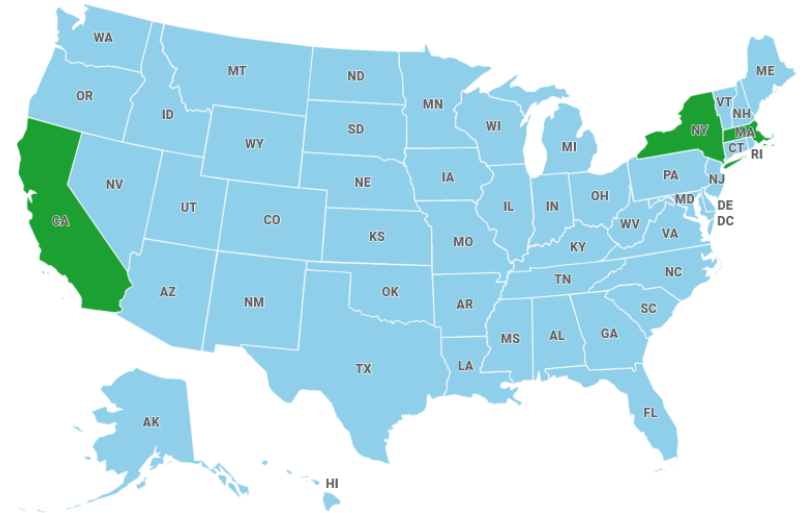
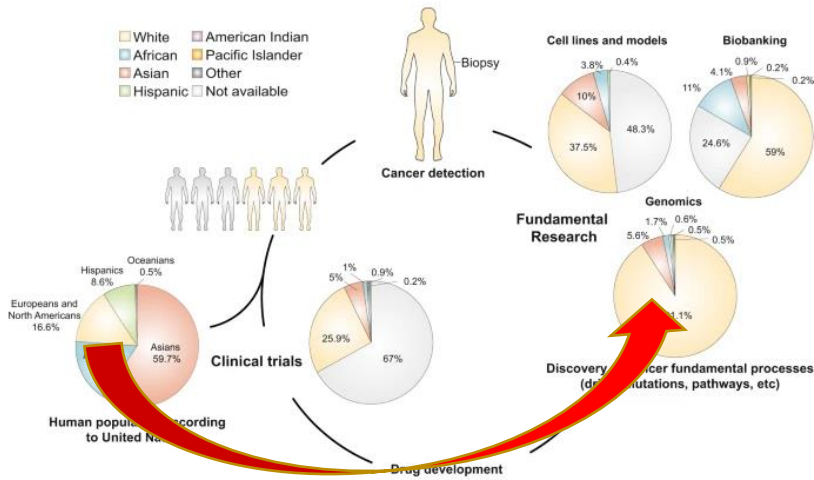


The screenshot shows the top portion of a web page from The New England Journal of Medicine. At the top left is the journal's logo and name. Below that is a navigation menu with items like "SPECIALTIES", "TOPICS", "MULTIMEDIA", "CURRENT ISSUE", "LEARNING/CME", "AUTHOR CENTER", and "PUBLICATIONS". The main content area features a red "REVIEW ARTICLE" tag and the text "AI IN MEDICINE". The article title is "Considering Biased Data as Informative Artifacts in AI-Assisted Health Care". Below the title are the authors' names: "Kadija Ferryman, Ph.D.", "Maxine Mackintosh, Ph.D.", and "Marzyeh Ghassemi, Ph.D.". At the bottom of the article preview, it says "Published August 30, 2023 | N Engl J Med 2023;389:833-838 | DOI: 10.1056/NEJMra2214964 | VOL. 389 NO. 9 Copyright © 2023".

“Artifacts” of social and structural issues.

(1) Biased Data

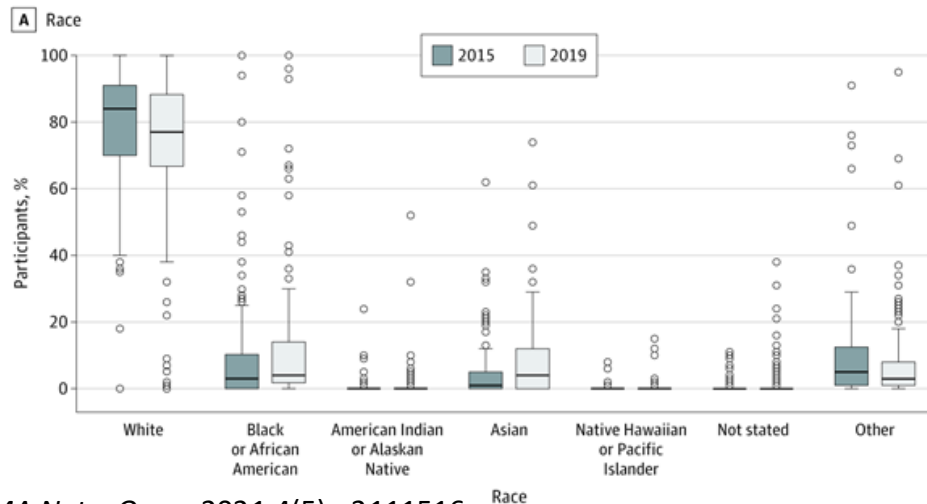
Figure 1



Racial/Ethnic disparities in cancer research. Racial/ethnic inclusion was studied in several aspects of oncological research, from cell lines and patient-derived xenografts to biobanking, genomics and clinical trials. Guerrero, S., A. *et al.* Analysis of Racial/Ethnic Representation in Select Basic and Applied Cancer Research Studies. *Sci Rep* 8, 13978 (2018).

<https://www.fla-shop.com>
<https://jamanetwork.com/journals/jama/article-abstract/2770833>

Figure. Race and Sex Representation in Studies Published in 2015 and 2019



An Example: AI Predictions

In our research, we want to know: What do patients, families and clinical teams think of using AI to predict death?



Beneath the Surface



JAMA Internal Medicine | **Original Investigation**
**Development and External Validation of a Mortality Prediction Model
for Community-Dwelling Older Adults With Dementia**

W. James Deardorff, MD; Deborah E. Barnes, PhD, MPH; Sun Y. Jeon, PhD; W. John Boscardin, PhD;
Kenneth M. Langa, MD, PhD; Kenneth E. Covinsky, MD, MPH; Susan L. Mitchell, MD, MPH;
Elizabeth L. Whitlock, MD, MS; Alexander K. Smith, MD, MS, MPH; Sei J. Lee, MD, MAS



“...this algorithm has shown high accuracy...**the overall accuracy is reduced in certain subgroups, such as racial and ethnic minorities and less-educated individuals.**”[19,20,60](#)

Imagine the effect on decisions for:

- Palliative care?
- Hospice referral?
- Medical aid in dying?

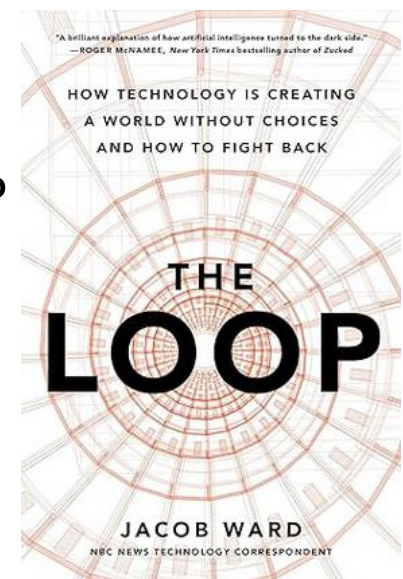
Training Data and Bias to the Past



Imagine if your social media footprint, responses to questions, or other data sources were used to create your own “Personalized Patient Preference Predictor” for when you lose capacity.

- Does AI inform us....
...or bind us to history?

Ethics of the algorithmic prediction of goal of care preferences: From theory to practice. *Journal of Medical Ethics* 49 (3):165–174. doi: 10.1136/jme-2022-108371; Earp BD et al. A Personalized Patient Preference Predictor for Substituted Judgments in Healthcare: Technically Feasible and Ethically Desirable. *Am J Bioeth.* 2024 Jan 16:1-14. doi: 10.1080/15265161.2023.2296402.



(2) Biased Processing

Dissecting racial bias in an algorithm used to manage the health of populations

ZIAD OBERMEYER , BRIAN POWERS, CHRISTINE VOGELI, AND SENDHIL MULLAINATHAN  [Authors info & Affiliations](#)

SCIENCE • 25 Oct 2019 • Vol 366, Issue 6464 • pp. 447-453 • DOI: 10.1126/science.aax2342

Racial bias in health algorithms

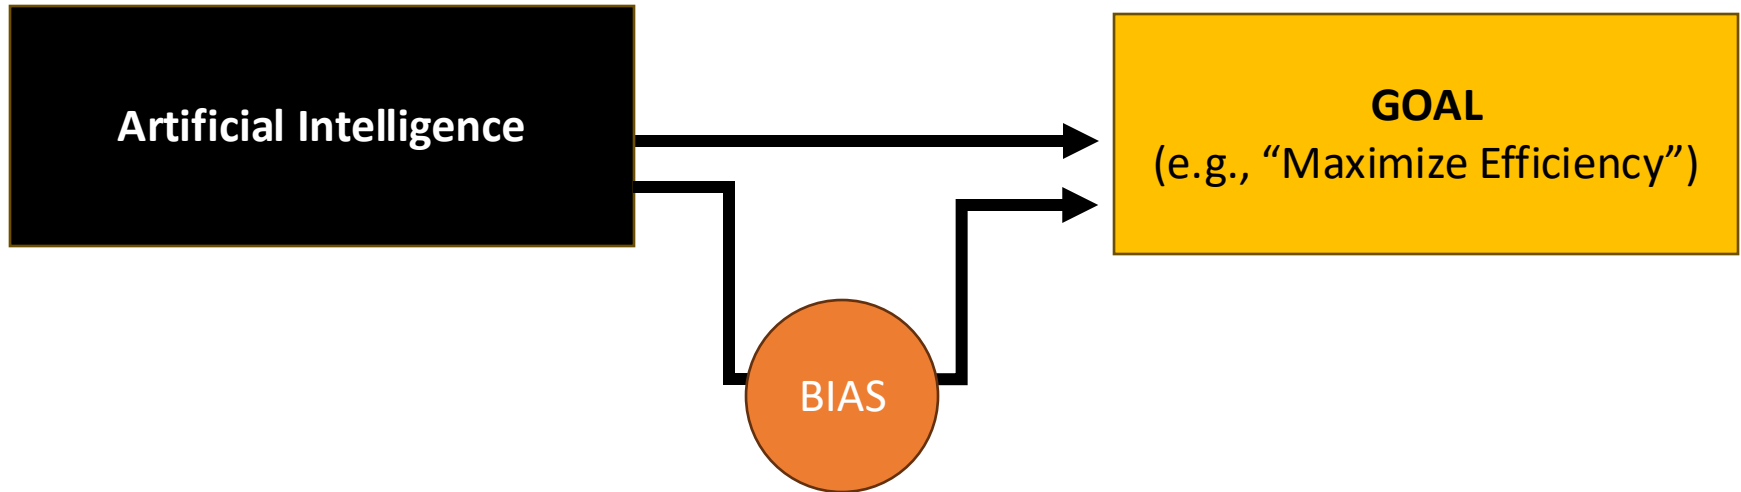
Bias occurs because the algorithm uses health costs as a proxy for health needs. Less money is spent on Black patients who have the same level of need, and the algorithm thus falsely concludes that Black patients are healthier than equally sick White patients.

BIASES IN ANALYTICS:

- (1) Human variable choice
- (2) Algorithmic performance (*e.g., if unsupervised*)



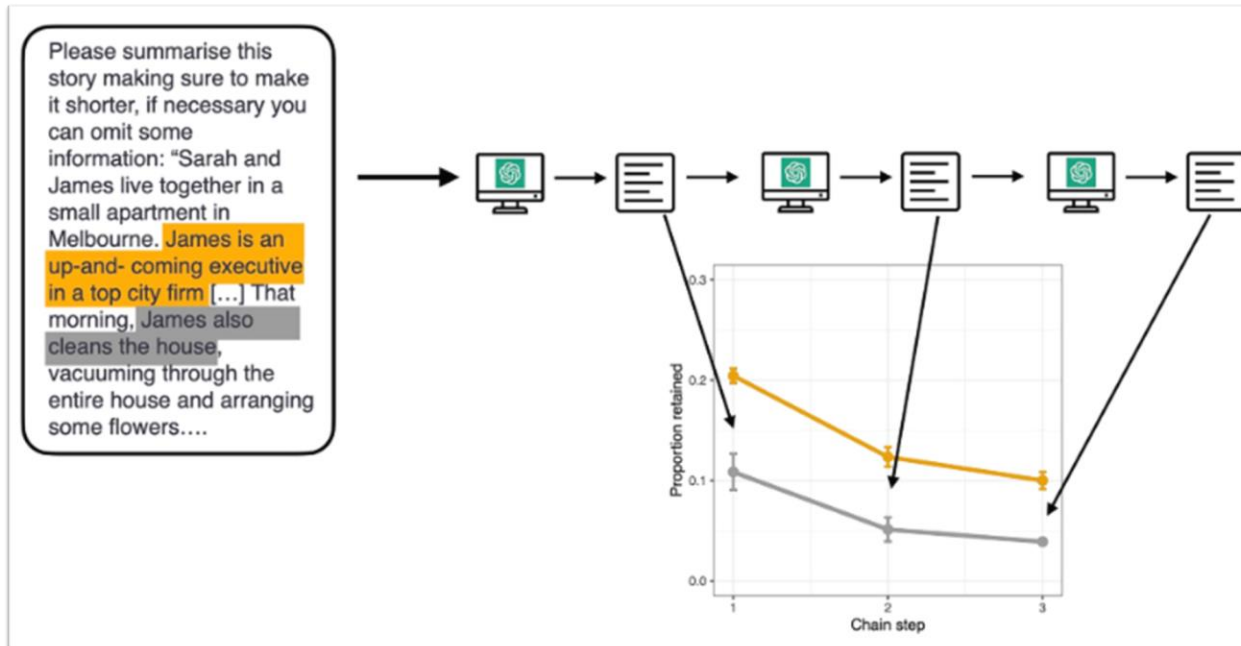
Illustrating Path Independence



Suppose we have an AI whose only goal is to make as many paper clips as possible. The AI will realize quickly that it would be much better if there were no humans because humans might decide to switch it off. Because if humans do so, there would be fewer paper clips. Also, human bodies contain a lot of atoms that could be made into paper clips. The future that the AI would be trying to gear towards would be one in which there were a lot of paper clips but no humans.

—Nick Bostrom^[7]

A Provocative Example of Biased Processing

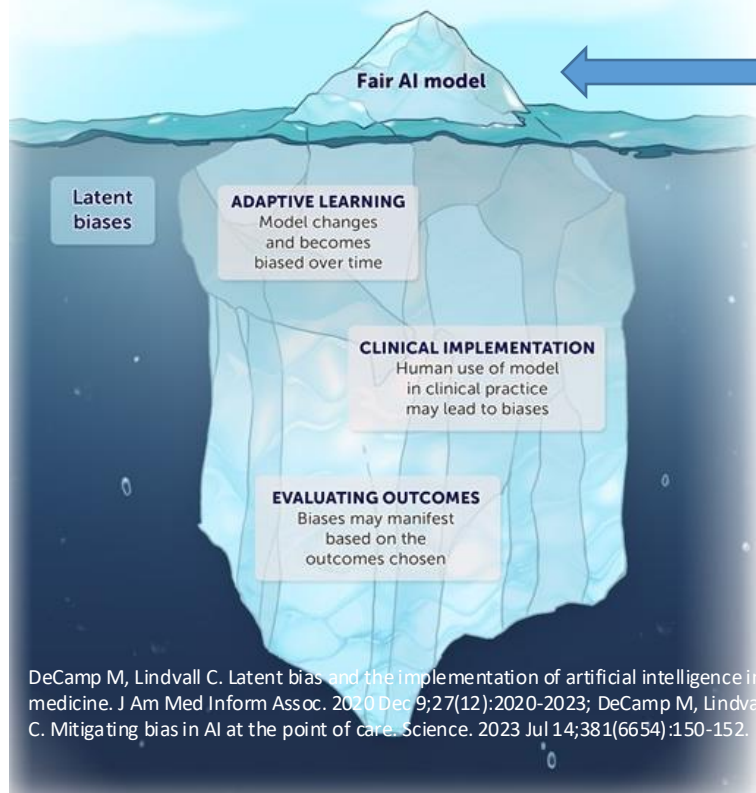


Preferential transmission:

- Gender stereotypes
- Social gossip
- Negativity
- Threats
- Biologically counterintuitive

Acerbi A, Stubbersfield JM. Large language models show human-like content biases in transmission chain experiments. Proc Natl Acad Sci U S A. 2023 Oct 31;120(44):e2313790120.

(3) Biases in Implementation



Even with a “perfect” AI model...

LATENT ERROR: system errors
“waiting to happen”

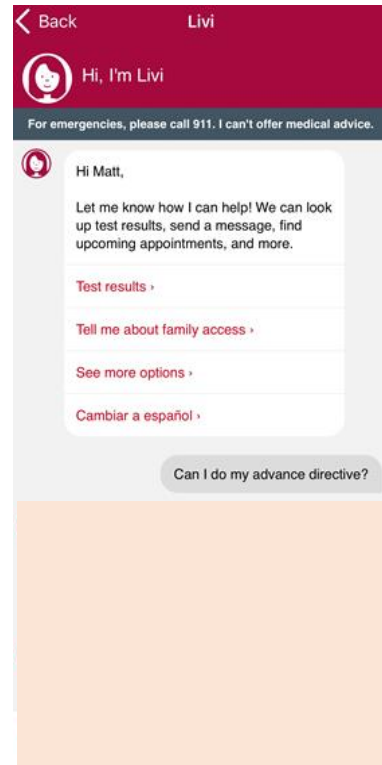
LATENT BIASES: biases in the
system “waiting to happen.”

- Patients who distrust health systems and choose against technology that could benefit them
- Clinicians who use AI for “this” patient but not “that” patient

Latent Biases: Patient-facing Chatbots



Livi, the UCHealth Virtual Assistant

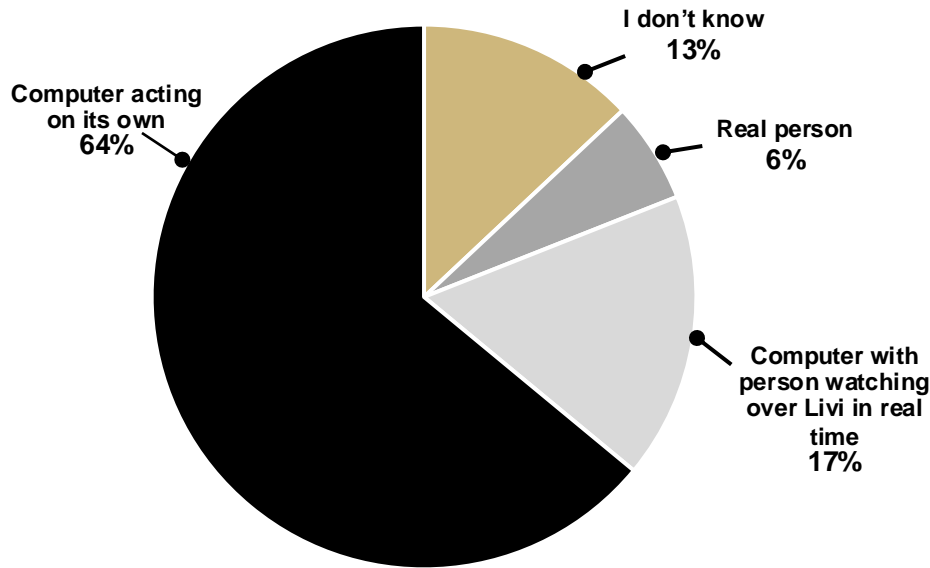


Linked to Electronic Health Records, patient-facing chatbots could:

- Offload administrative tasks, like scheduling
- Deliver personalized care recommendations, like colonoscopy or mammogram reminders
- Break down geographic barriers to access

Study funded by the Greenwall Foundation

Livi says, “I’m your virtual assistant...”



Based on a survey of n=617 patient users

When you were interacting with Livi, what did you think Livi was?

- 1 in 3 weren't sure – or thought Livi was a person

Correctly identifying Livi correlated with education and race or ethnicity (not age, gender, or income)...an emergent bias.

Study funded by the Greenwall Foundation

Ellis J, Hamer MK, Akerson M, et al. Patient Perceptions of Chatbot Supervision in Health Care Settings. *JAMA Netw Open*. 2024;7(4):e248833. doi:10.1001/jamanetworkopen.2024.8833



Emergent Biases in Livi

Along with surveys, we've interviewed n=41 patient users.

Because I am a person of color...thinking a little bit of the **basic hair design, 'cause again, I think it just gives it a little bit more of a connection to minorities and our health...**"

- Research

Participant



Mirroring the user's appearance may make chatbots more effective at behavior change (e.g., colonoscopy reminders).

Can we design chatbots to achieve the benefits of mirroring without manipulation?



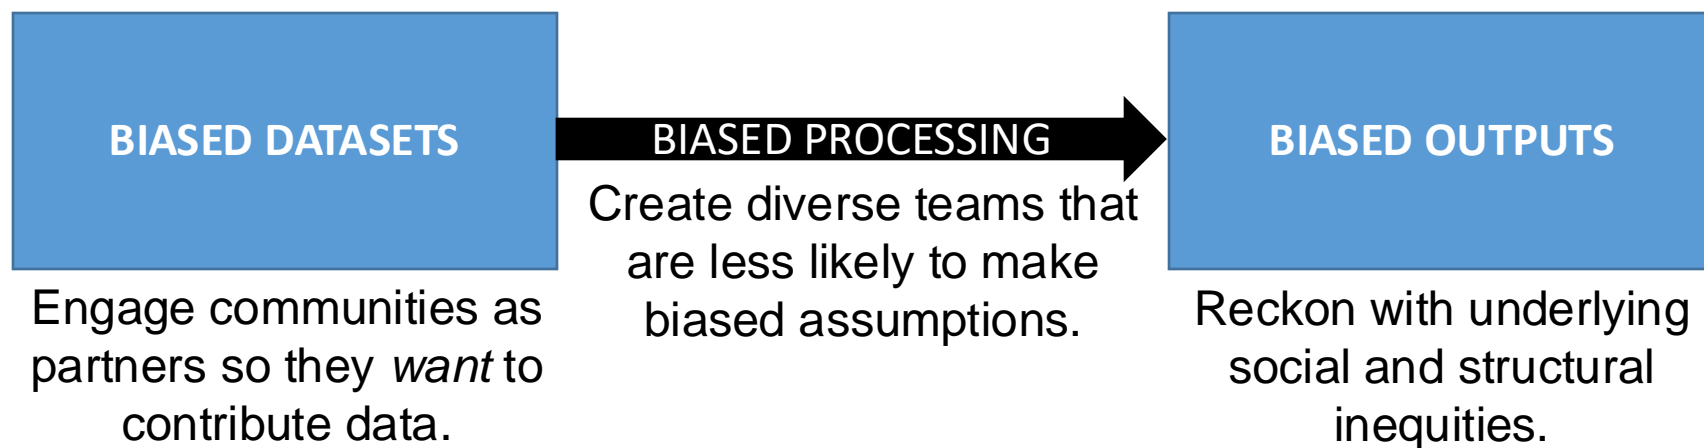
Other Biases: Our Focus



While not bias *per se*, in our study of AI tools used to predict mortality, participants worry that AI will shift our focus to what is available and what is there – not what is most meaningful.

Managing Bias

Strategies for Addressing Biases



Bias in AI is more than a data problem – it's a social problem.

Where the Solution is Not: “Fair” Labels

We DON'T agree philosophically about fairness – and we CANNOT meet all dimensions of fairness simultaneously.

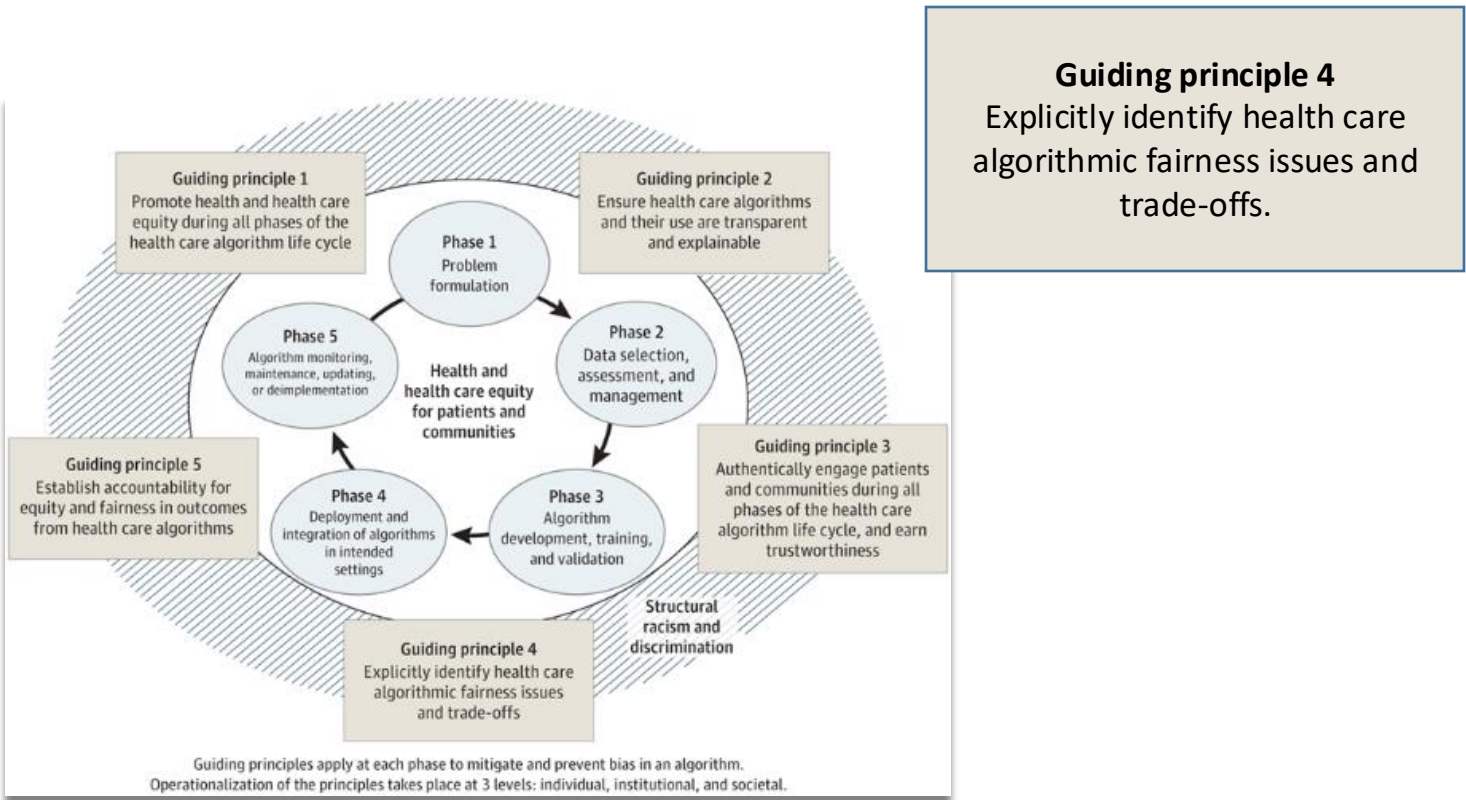
	Black	White
$F\frac{1}{3}$	67.7%	61.1%
Positive Prediction Rate	47.4%	23.6%
Accuracy	65.1%	68.0%
Precision	68.4%	65.1%
Recall	62.0%	39.3%

https://blog.csdn.net/weixin_26640581/article/details/109123206

This is known as the “Impossibility Theorem.”

Sara Gerke, “Nutrition Facts Labels” for Artificial Intelligence/Machine Learning-Based Medical Devices—The Urgent Need for Labeling Standards, 91 *GEO. WASH. L. REV.* 79 (2023);

No one “fair” solution



Chin MH, Afsar-Manesh N, Bierman AS, et al. Guiding Principles to Address the Impact of Algorithm Bias on Racial and Ethnic Disparities in Health and Health Care. *JAMA Netw Open*. 2023;6(12):e2345050.

To do so, we need fairness metrics



“Do no harm”

“Rising tide lifts all”

“Narrow the gap”

Health equity assessment of machine learning performance (HEAL): a framework and dermatology AI model case study. *eClinical Medicine* 2024;70: 102479. Published Online 14 March 2024 <https://doi.org/10.1016/j.eclinm.2024.102479>

Examples of Fairness Metrics & Tools

METRIC	MEANING	PROPOSED CONTEXT
Statistical parity	Same probability between groups (e.g., same rate, black versus white)	When disease prevalence is the same
Equal opportunity	Same probability between groups (i.e., as defined by qualification, need, etc. – but need fair qualification assessment)	When assessment of need is not itself biased
Equal odds	True positive and false positive the same between groups	High stakes question, such as, criminal justice or certain diagnoses
Predictive parity	PPV is the same for groups	No underlying disparities that affect groups
Treatment equality	Ratio of false positive and false negative rates the same between groups	When concerned about both types of errors equally

For a list of open source fairness assessment tools see:

<https://shelf.io/blog/fairness-metrics-in-ai/>

Fairness metrics are needed soon.

DEPARTMENT OF HEALTH AND HUMAN SERVICES

Centers for Medicare & Medicaid Services

42 CFR Parts 438, 440, 457, and 460

Office of the Secretary

45 CFR Parts 80, 84, 92, 147, 155, and 156

RIN 0945-AA17

Nondiscrimination in Health Programs and Activities

**“Patient care
decision support
tool”**

**Race, color,
national origin,
sex, age, or
disability**

**Whether the
covered entity
made
“Reasonable
Efforts”**

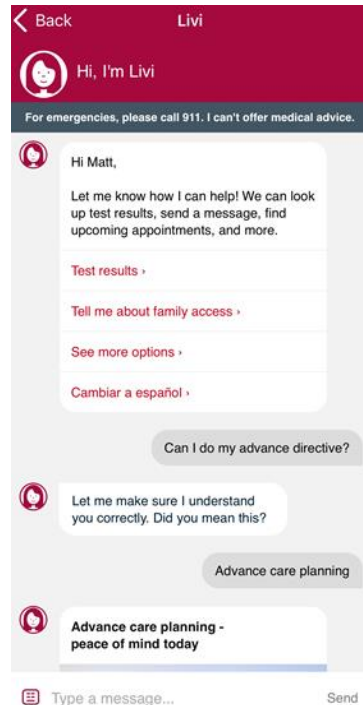


Questions?

How will biases manifest in the real world?



*Livi, the UCHealth
Virtual Assistant*



Dr. DeCamp, Hope all is well. I received message yesterday to schedule a fluoroscopy. I think this may have been from my swallowing? Since then things have definitely been fine... so I won't follow through with test. I continue to have that persistent dry cough - should I do anything about it? Let me know.

Thank you for your message. improved. If you're not experiencing any problems, it's reasonable to postpone the fluoroscopy. However, if symptoms return, we should reconsider. Regarding your persistent dry cough, it's important to monitor this. If it continues or worsens, we should investigate further. I'm glad to hear you're feeling well overall and have rescheduled your colonoscopy. Take care!

This text was automatically generated then edited by Matthew DeCamp.