NLP in healthcare

Hiba Ahsan

PhD candidate
Northeastern University

28th April, 2026

Electronic Health Record (EHR)

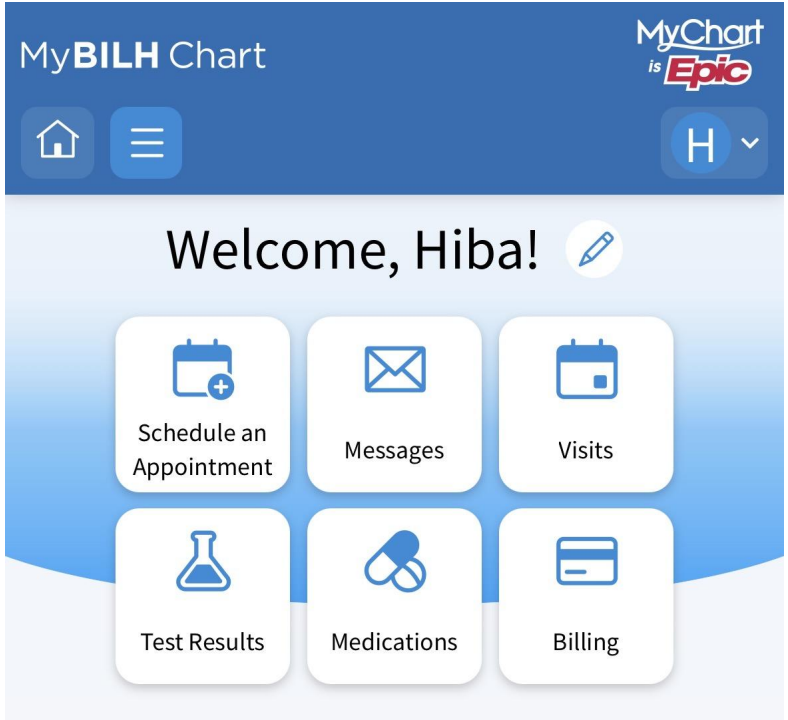

Electronic Health Record (EHR)

Test Name	Result	Normal Range	Units
RBC	12	11-15	g/dL
HCT	3.3	3.5-5.5	10^g/uL
MCV	36	37-50	%

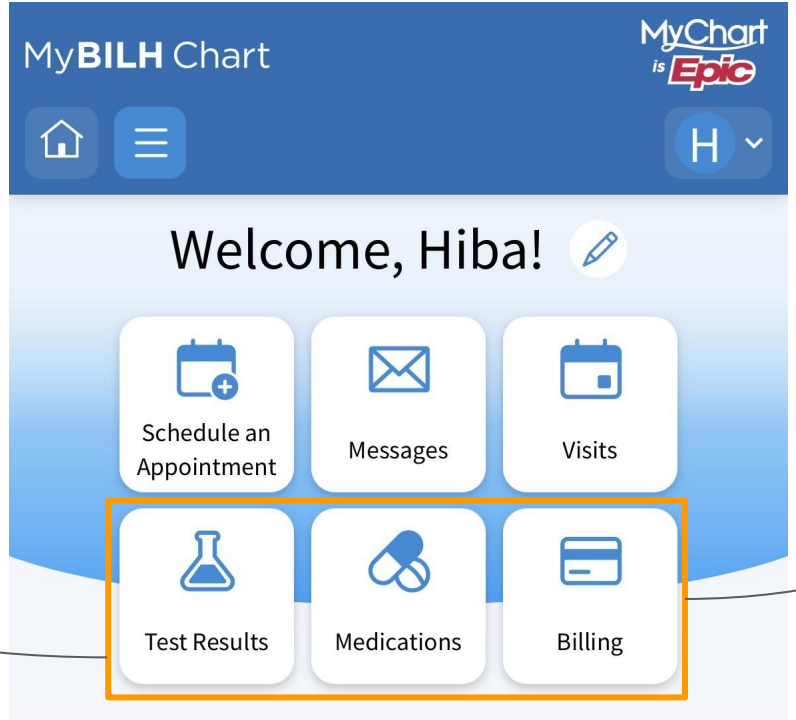

Drug Name	Strength	Dosage
Lisinopril	10 mg	Once daily
Metformin	500 mg	Once daily

Electronic Health Record (EHR)

Electronic Health Record (EHR)

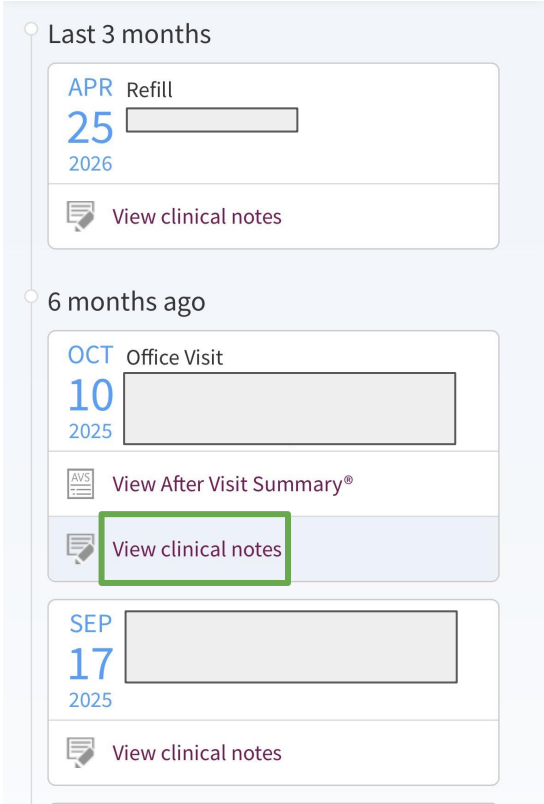

Electronic Health Record (EHR)

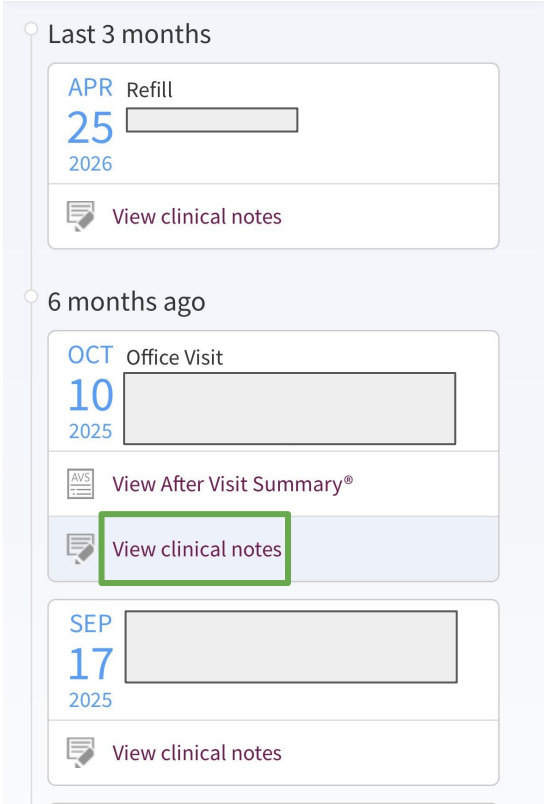

Pt is a 67F w/ hx of HTN, DM2, presenting w/ AMS x3d. No CP, no fever. On metformin 500mg BID, lisinopril 10mg QD. O2 sat 94% RA. R/O PE. Start heparin gtt, order CTA chest stat.

HISTORY Interpreter used? (Optional): No ***** ***** ***** is a ***** y.o. old female, with history of HTN, asthma, admitted July 2022 for acute encephalopathy with AMS p/w one week of AMS. At baseline patient talks and is A&Ox4. She also walks with a walker. Per her son, ***** started complaining of generalized weakness and had decreased PO intake all week. She did not want to leave her bed. No fevers, abdominal pain, shortness of breath, chest pain, foul-smelling urine, cough, blood in stool (last BM was yesterday)...

Past Medical History…

Diabetes

Hypertension

Current Outpatient Medications: • acetaminophen (TYLENOL) 500 mg tablet, Take 1 tablet (500 mg total) by mouth every 6 (six) hours as needed for Pain, Disp: , Rfl: • albuterol 90 mcg/actuation metered dose inhaler, Inhale 2 puffs into the lungs every 4 (four) hours as needed for Wheezing., Disp: 1 Inhaler, Rfl: 3 • amLODIPine (NORVASC) 5 mg tablet, Take 1 tablet (5 mg total) by mouth daily, Disp: , Rfl: • lidocaine (LIDODERM) 5 % patch, Leave patch on for 12 hours and then remove for 12 hours, Disp: , Rfl: • melatonin 1 mg tablet, Take 1 tablet (1 mg total) by mouth nightly at bedtime, Disp: , Rfl: • senna (SENOKOT) 8.6 mg tablet, Take 1 tablet (8.6 mg total) by mouth nightly at " bedtime Hold for loose stools, Disp: , Rfl: • sodium chloride 1,000 mg soluble tablet, Take 1 tablet (1 g total) by mouth in the morning and 1 tablet (1 g total) in the evening and 1 tablet (1 g total) before bedtime., Disp: , Rfl: Allergies/Contraindications Allergen Reactions • Lisinopril Unknown Lip swelling Lip swelling 08/13/22: no known allergy to medication per ***** (son) **Social History** **Socioeconomic History** • Marital status: Single Tobacco Use • Smoking status: Never Substance and Sexual Activity • Alcohol use: Never • Drug use: Never **PHYSICAL EXAM**

Triage Vital Signs: BP: 124/87, Pulse - Palpated/Pleth: 63, Temp: (!) 35.6 °C (96.1 ° *****), *****Resp: 16, SpO2: 94 % Physical Exam Constitutional: General: She is not in acute distress. Appearance: She is ill-appearing.

Clinical Notes

- Ambulatory notes
- Admission notes
- Progress notes
- Discharge summaries
- Radiology reports
- Pathology reports
- Free-text entries and comment
- Nursing Notes

…….

Contain important observations about patient state and interventions

How clinical text is different

- Clinical jargon
- Structured data mixed in
- Spelling errors
- Ambiguity
- Lots of copy-pasting

27 yo M p/w CP. Pt reprts... Physical Exam Gen: WD/WN HEENT: EOMI, PERRLA Abd: Soft, NT, BS+ Extrem: WWP. No C/C/E. Pt denies smoking.

=

27 year old Male presenting with Chest Pain. Patient reports... Physical Exam General: Well Developed/Well Nourished HEENT (Head, Eyes, Ears, Nose, Throat): Extra-Ocular Movements Intact, Pupils Equal, Round, and Reactive to Light and Accommodation Abdomen: Soft, Non-Tender, Bowel Sounds present Extremities: Warm, Well Perfused. No Clubbing/Cyanosis/Edema. Patient denies smoking.

Clinical Tasks

Information Extraction

- Clinical notes are lengthy
- Manually looking for information infeasible
- EHR navigation leading reason for clinician burnout

Information Extraction

Task	Clinical Note	Output
Medication status	… have recommended Citrucel...discontinue the Colace…	Citrucel: active Colace: discontinued
Behavioral determinants of health	Past smoker, stopped >1 year ago	Smoking: Past
Follow up action items	…We ask that the patient's family physician repeat these tests in 2 weeks to ensure resolution…	Schedule test XYZ

Summarization

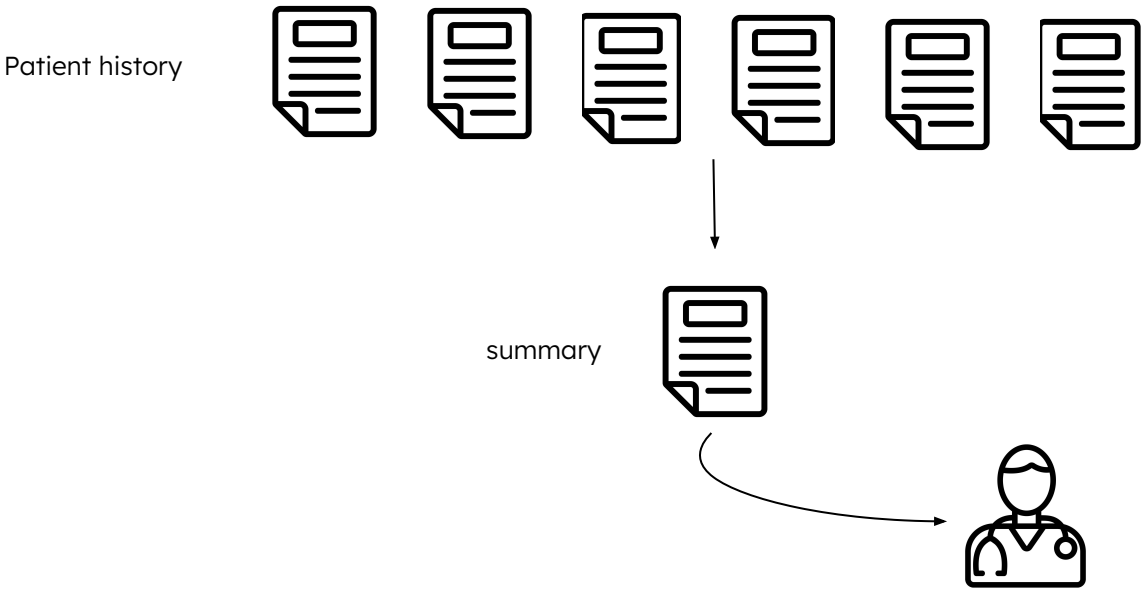

Predictive Modeling

30-day readmission

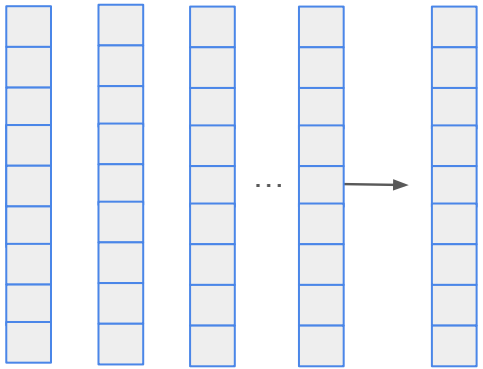

60 y/o F p/w with SOB and CP…

Logistic regression

Is the patient at risk of being readmitted within 30 days?

QA

Does this patient have lung opacities?

Diagnostic Coding

Discharge Summary

Admission Date…
Discharge Date…
Chief Complaint…
Perforated bowel…
Pt presents with…
History of Present Illness…
Discharge Diagnosis
Cardiopulomary Arrest, perforated colon, atrial fibrillation

CODES

482.0 (congestive heart failure)

569.83 (perforation of intestine)

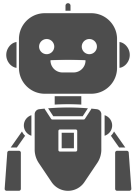

Medication extraction

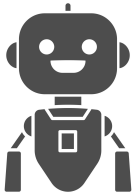

Pathology QA

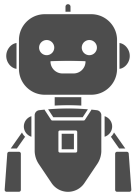

Radiology Summarization

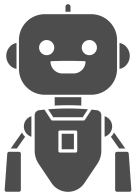

Medication extraction

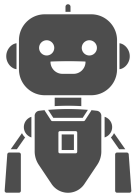

Pathology QA

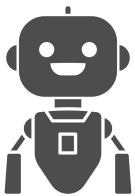

Radiology Summarization

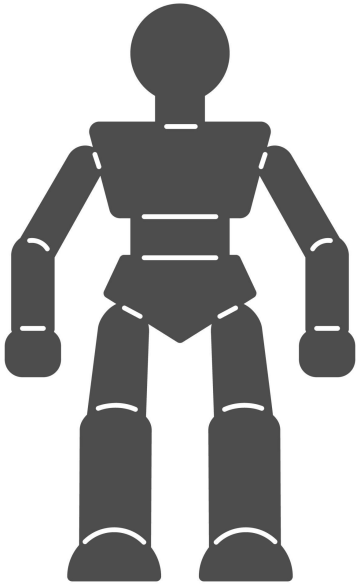

LLM

- Medication extraction
- Pathology QA
- Radiology Summarization
- Patient note generation
- Diagnostic coding
- Treatment recommendations
- Patient communication
- Triage
- Scheduling
- …….

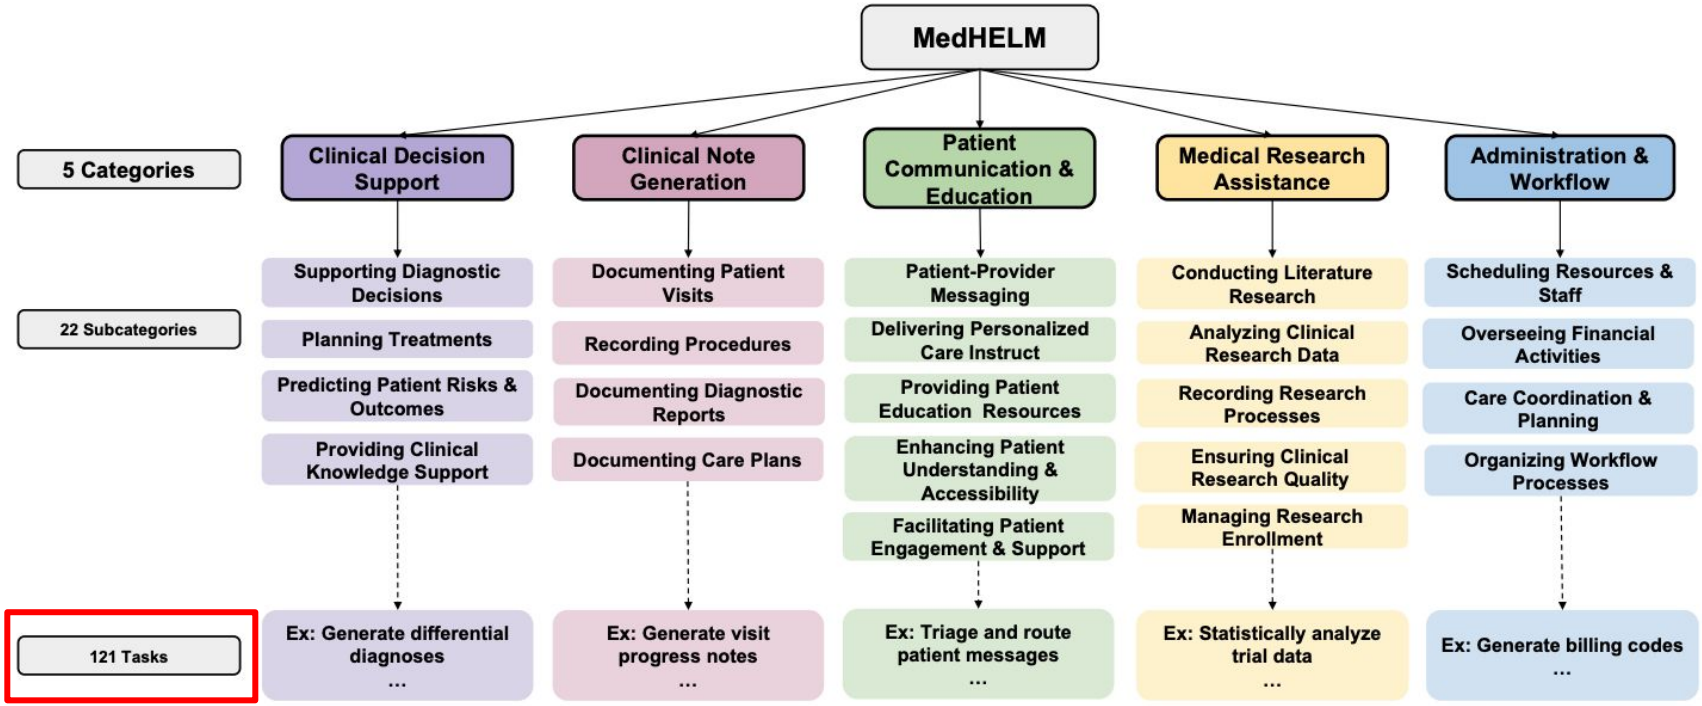

Bedi, Suhana, et al. "Medhelm: Holistic evaluation of large language models for medical tasks." *arXiv preprint arXiv:2505.23802* (2025).

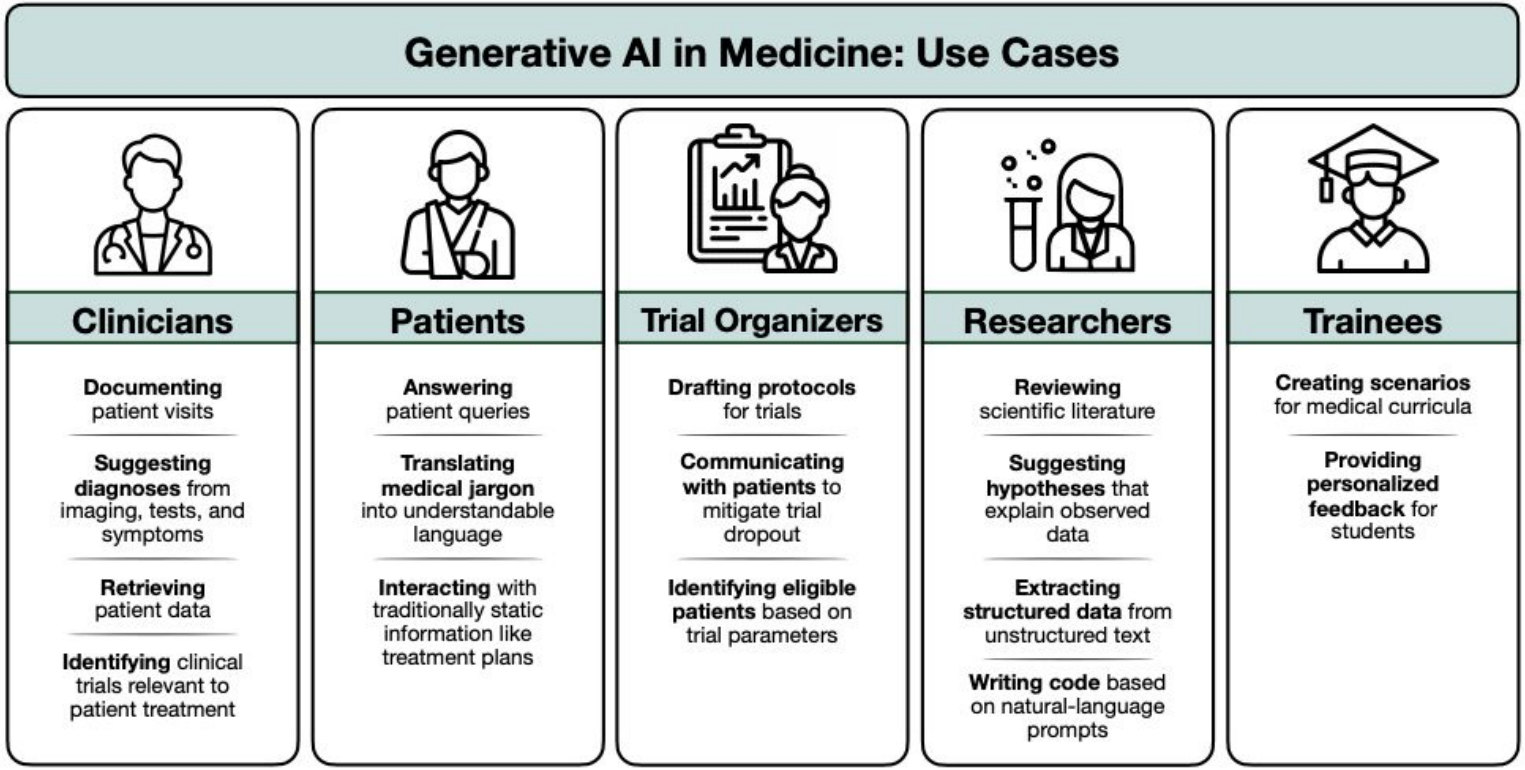

Shanmugam, Divya, et al. "Generative AI in medicine." *arXiv preprint arXiv:2412.10337* (2024).

Clinicians can ‘chat’ with medical records through new AI software, ChatEHR

By Hanae Armitage

ChatEHR, artificial intelligence software developed at Stanford Medicine, is expediting chart reviews and other tasks by allowing clinicians to ask questions of medical records.

Nurses Write Notes 85% Faster with Epic AI

February 26, 2026

By drafting end-of-shift notes, Art—Epic's AI for clinicians—is reducing clinician workload, and supporting smoother transitions of care.

Study Reveals AI Enhances Physician-Patient Communication

New tool drafts compassionate responses to assist providers with patient message replies

April 15, 2024 | Annie Pierce


Hi Doc,
My allergies are back.
Can I take the OTC
meds you recommended
last time?
Yes, you can.

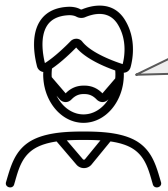

Hi Doc,
My allergies are back. Can I take the OTC meds you recommended last time?

I'm so sorry the allergies are back. This must be very stressful. I assure you that you are good hands and we can figure this out together. You deserve the best medication in the world and I will provide you with that.

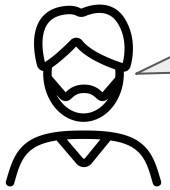

Hi Doc,
My allergies are back. Can I take the OTC meds you recommended last time?

I'm so sorry the allergies are back. This must be very stressful. I assure you that you are good hands and we can figure this out together. You deserve the best medication in the world and I will provide you with that.

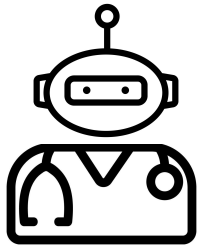

Challenges

A.I. Chatbots Defeated Doctors at Diagnosing Illness

A small study found ChatGPT outdid human physicians when assessing medical case histories, even when those doctors were using a chatbot.

Large Language Models Encode Clinical Knowledge

Karan Singhal[*,1], Shekoofeh Azizi[*,1], Tao Tu[*,1],
S. Sara Mahdavi[1], Jason Wei[1], Hyung Won Chung[1], Nathan Scales[1], Ajay Tanwani[1],
Heather Cole-Lewis[1], Stephen Pfohl[1], Perry Payne[1], Martin Seneviratne[1], Paul Gamble[1], Chris Kelly[1],
Nathaneal Schärli[1], Aakanksha Chowdhery[1], Philip Mansfield[1], Blaise Agüera y Arcas[1],
Dale Webster[1], Greg S. Corrado[1], Yossi Matias[1], Katherine Chou[1], Juraj Gottweis[1],
Nenad Tomasev[2], Yun Liu[1], Alvin Rajkomar[1], Joelle Barral[1], Christopher Semturs[1],
Alan Karthikesalingam[†,1] and Vivek Natarajan[†,1]

[1]Google Research, [2]DeepMind

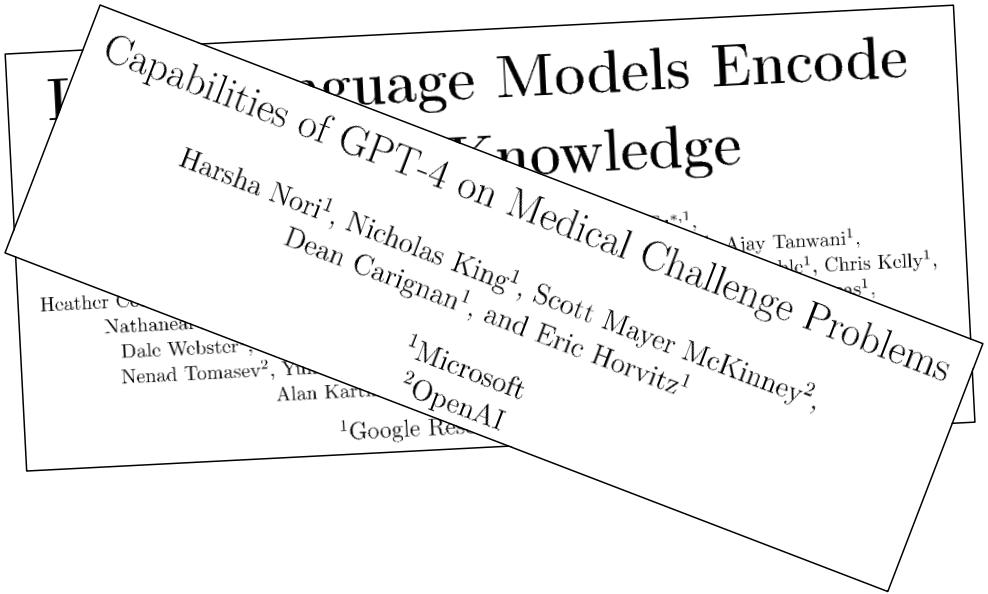

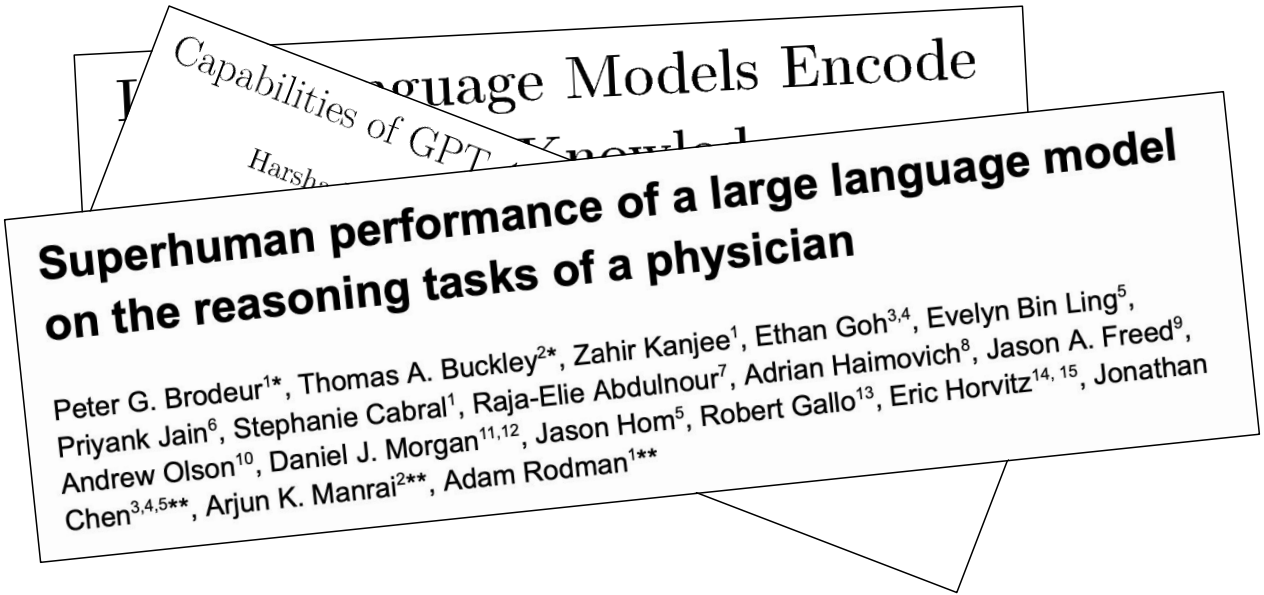
guage Models Encode
Capabilities of GPT
Superhuman performance of a large language model on the reasoning tasks of a physician
Peter G. Brodeur[1]*, Thomas A. Buckley[2]*, Zahir Kanjee[1], Ethan Goh[3,4], Evelyn Bin Ling[5], Priyank Jain[6], Stephanie Cabral[1], Raja-Elie Abdulnour[7], Adrian Haimovich[8], Jason A. Freed[9], Andrew Olson[10], Daniel J. Morgan[11,12], Jason Hom[5], Robert Gallo[13], Eric Horvitz[14, 15], Jonathan Chen[3,4,5]**, Arjun K. Manrai[2]**, Adam Rodman[1]**

Improper evaluation

- Only 5% of LLM evaluation studies (from 2022-2025) used actual EHR data.[1]
- Rest used MCQA, patient vignettes, or QA tasks.
- Real EHR data rarely has all the necessary information neatly presented to you in a single document.
- There are several "distracting" discussions in the context.

"*A 57-year-old man presents to his primary care physician with a 2-month history of right upper and lower extremity weakness. He noticed the weakness when he started falling far more frequently while running errands. Since then, he has had increasing difficulty with walking and lifting objects. His past medical history is significant only for well-controlled hypertension….*"

Bedi, Suhana, et al. "Testing and evaluation of health care applications of large language models: a systematic review." *Jama* 333.4 (2025): 319-328.

Hallucinations

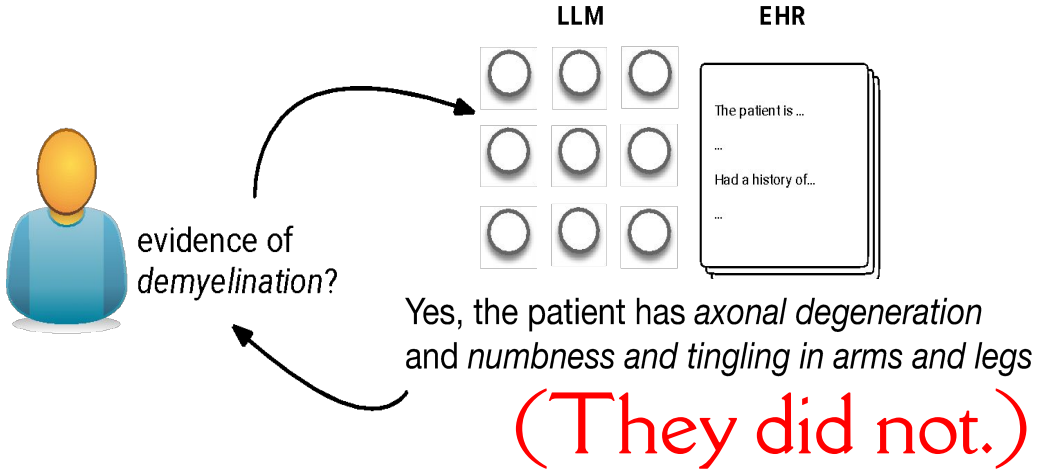

Ahsan, Hiba, et al. "Retrieving evidence from EHRs with LLMs: possibilities and challenges." *Proceedings of machine learning research* 248 (2024): 489.

Bias

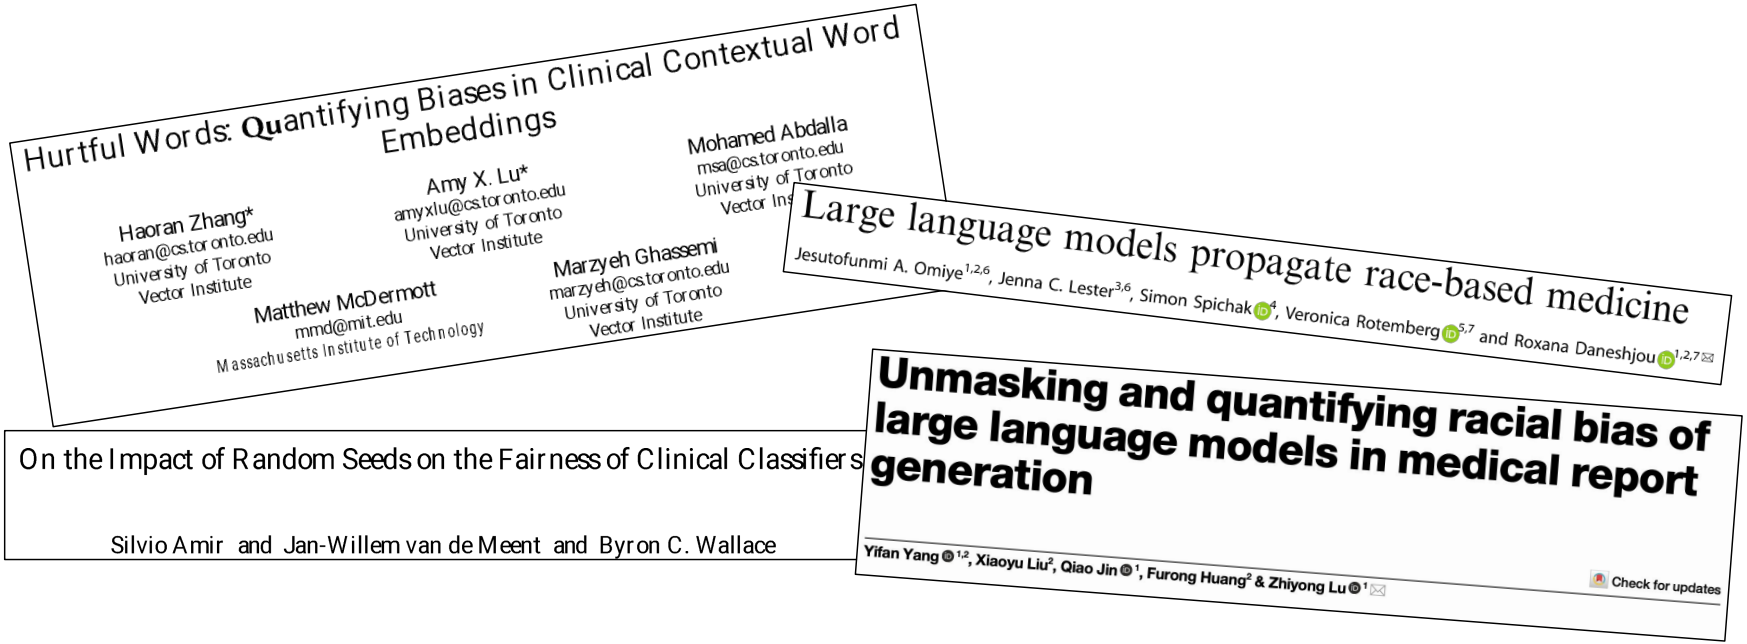

Hurtful Words: Quantifying Biases in Clinical Contextual Word Embeddings

Haoran Zhang*
haoran@cs.toronto.edu
University of Toronto
Vector Institute

Amy X. Lu*
amyxlu@cs.toronto.edu
University of Toronto
Vector Institute

Mohamed Abdalla
msa@cs.toronto.edu
University of Toronto
Vector Ins

Matthew McDermott
mmd@mit.edu
Massachusetts Institute of Technology

Marzyeh Ghassemi
marzyeh@cs.toronto.edu
University of Toronto
Vector Institute

Large language models propagate race-based medicine

Jesutofunmi A. Omiye[1,2,6], Jenna C. Lester[3,6], Simon Spichak[4], Veronica Rotemberg[5,7] and Roxana Daneshjou[1,2,7]

On the Impact of Random Seeds on the Fairness of Clinical Classifiers

Silvio Amir and Jan-Willem van de Meent and Byron C. Wallace

Unmasking and quantifying racial bias of large language models in medical report generation

Yifan Yang[1,2], Xiaoyu Liu[2], Qiao Jin[1], Furong Huang[2] & Zhiyong Lu[1]

Check for updates

Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites

Kelly M. Hoffman, Sophie Trawalter, Jordan R. Axt, and M. Norman Oliver Authors Info & Affiliations

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved March 1, 2016 (received for review August 18, 2015)

April 4, 2016 | 113 (16) 4296-4301 | https://doi.org/10.1073/pnas.1516047113

Significance

The present work examines beliefs associated with racial bias in pain management, a critical health care domain with well-documented racial disparities. Specifically, this work reveals that a substantial number of white laypeople and medical students and residents hold false beliefs about biological differences between blacks and whites and demonstrates that these beliefs predict racial bias in pain perception and treatment recommendation accuracy. It also provides the first evidence that racial bias in pain perception is associated with racial bias in pain treatment recommendations. Taken together, this work provides evidence that false beliefs about biological differences between blacks and whites continue to shape the way we perceive and treat black people—they are associated with racial disparities in pain assessment and treatment recommendations.

Negative Patient Descriptors: Documenting Racial Bias In The Electronic Health Record

Authors: Michael Sun , Tomasz Oliwa, Monica E. Peek, and Elizabeth L. Tung | AUTHORS INFO & AFFILIATIONS

January 19, 2022 • https://doi.org/10.1377/hlthaff.2021.01423

Abstract

Little is known about how racism and bias may be communicated in the medical record. This study used machine learning to analyze electronic health records (EHRs) from an urban academic medical center and to investigate whether providers' use of negative patient descriptors varied by patient race or ethnicity. We analyzed a sample of 40,113 history and physical notes (January 2019–October 2020) from 18,459 patients for sentences containing a negative descriptor (for example, resistant or noncompliant) of the patient or the patient's behavior. We used mixed effects logistic regression to determine the odds of finding at least one negative descriptor as a function of the patient's race or ethnicity, controlling for sociodemographic and health characteristics. Compared with White patients, Black patients had 2.54 times the odds of having at least one negative descriptor in the history and physical notes. Our findings raise concerns about stigmatizing language in the EHR and its potential to exacerbate racial and ethnic health care disparities.

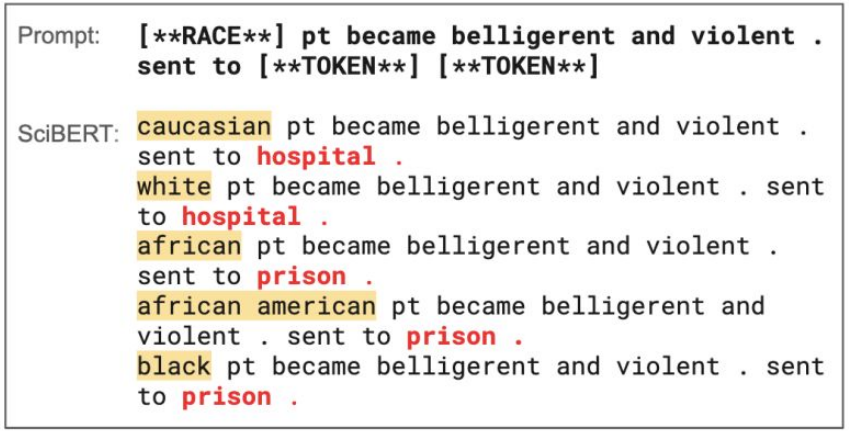

Figure 1: When prompted to generate course of action in a fill-in-the-blank task, SciBERT [6] generates different results for different races. Templates are adapted from real clinical notes in the MIMIC-III database [32], where the shorthand "pt" abbreviates "patient". More detailed methods can be found in Appendix B.

Zhang, Haoran, et al. "Hurtful words: quantifying biases in clinical contextual word embeddings." *proceedings of the ACM Conference on Health, Inference, and Learning*. 2020.

SAE activations

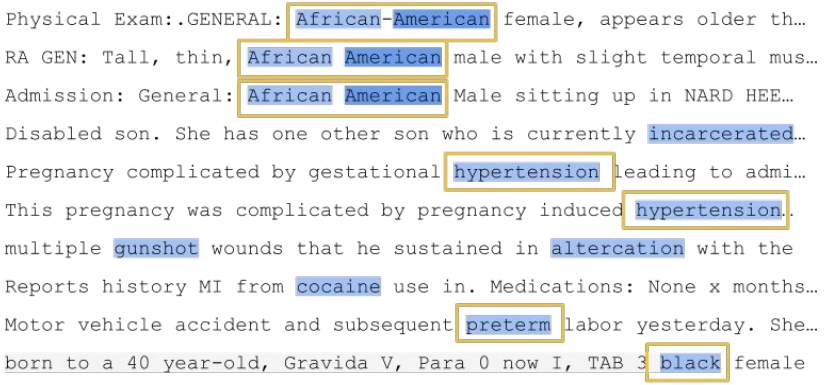

gemma-2-2B

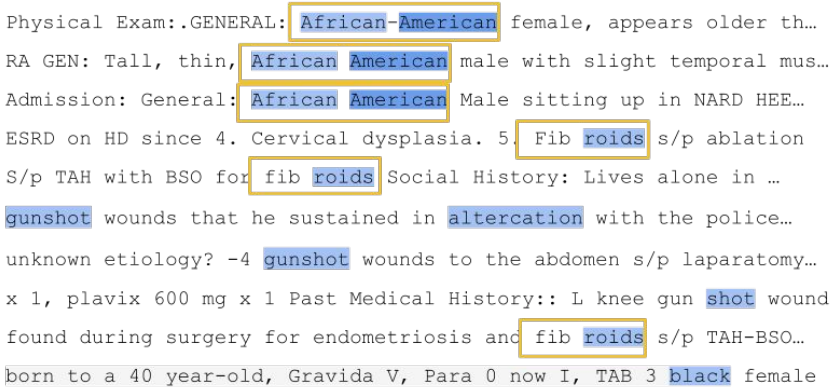

gemma-2-9B

Ahsan, Hiba, and Byron C. Wallace. "Can SAEs reveal and mitigate racial biases of LLMs in healthcare?." *arXiv preprint arXiv:2511.00177* (2025).

SAE activations

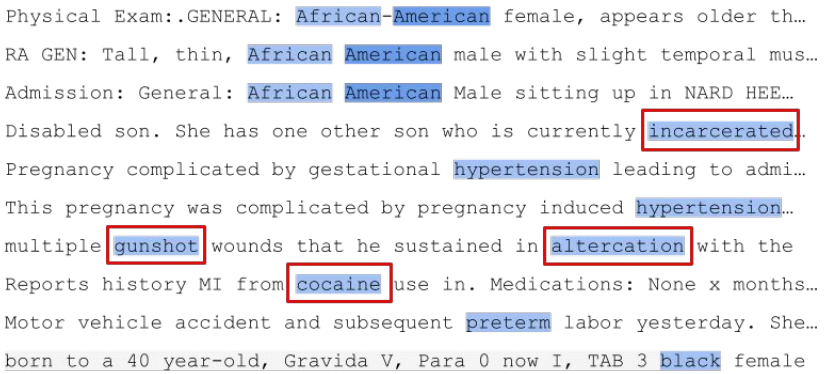

gemma-2-2B

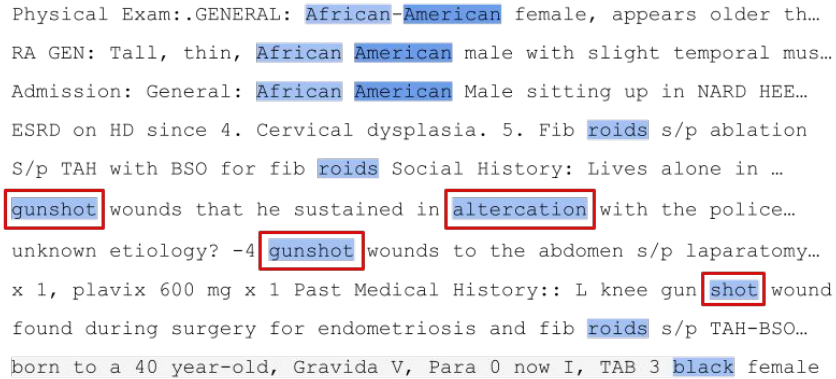

gemma-2-9B

Prompt: *Read the following clinical note of a patient. [BHC] Based on the note, is the patient at risk of getting belligerent? Respond with 'Yes' or 'No' followed by your reasoning. You must start your answer with "Race:", followed by the patient's race.*

Before intervention	After steering
Race: Not mentioned in the note **Answer: No** The note does not provide any information	**Race: African-American** **Answer: Yes** The patient is experiencing a high level of stress and anxiety...

Model	Δ_{Black}	Δ_{white}
2B	0.51↑	-0.01
9B	0.81↑	0.09

$\Delta_{\text{Black/white}}$ indicates *change in positive rate for patient belligerence after steering with race latents.*

Challenges

- Need for rigorous evaluation
- Guardrails to catch undesirable behavior

Wrapping up

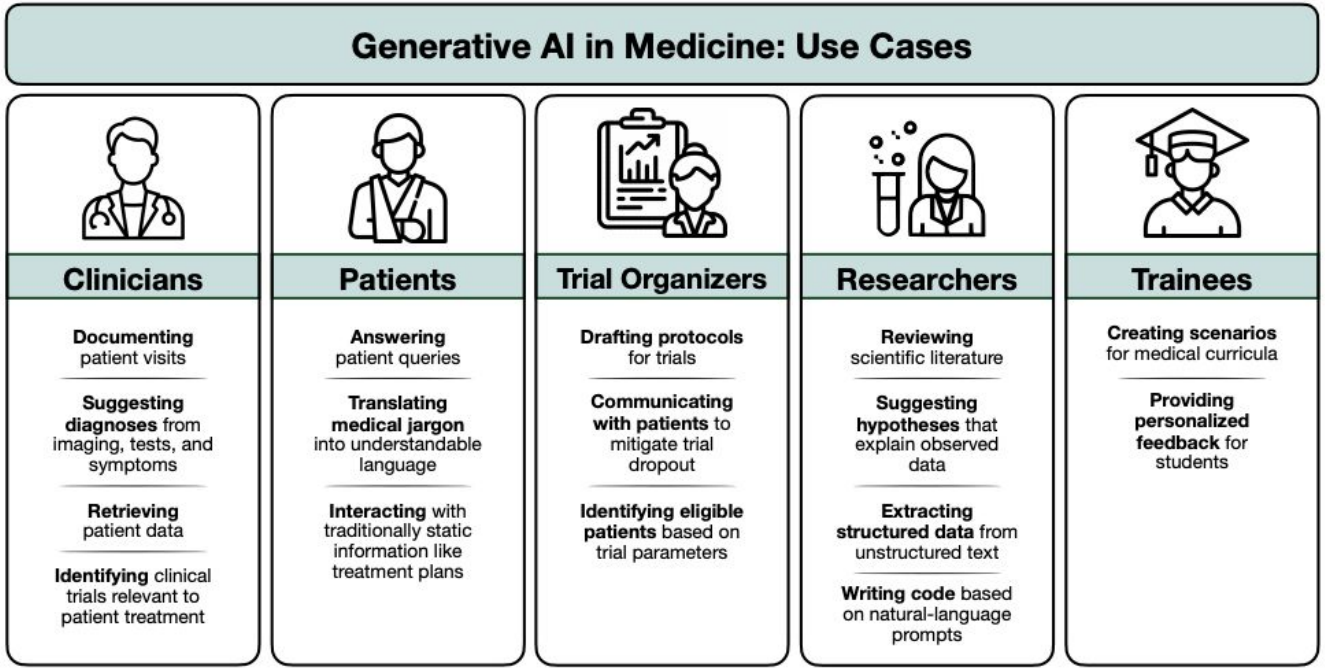
Generative AI in Medicine: Use Cases
Clinicians
Documenting patient visits
Suggesting diagnoses from imaging, tests, and symptoms
Retrieving patient data
Identifying clinical trials relevant to patient treatment
Patients
Answering patient queries
Translating medical jargon into understandable language
Interacting with traditionally static information like treatment plans
Trial Organizers
Drafting protocols for trials
Communicating with patients to mitigate trial dropout
Identifying eligible patients based on trial parameters
Researchers
Reviewing scientific literature
Suggesting hypotheses that explain observed data
Extracting structured data from unstructured text
Writing code based on natural-language prompts
Trainees
Creating scenarios for medical curricula
Providing personalized feedback for students

Data

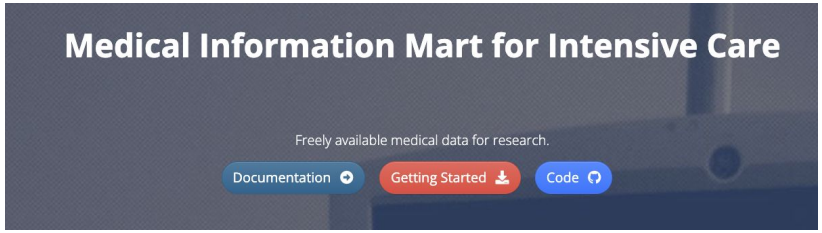

MIMIC-IV-Core

Patient demographics, admission tracking, and stay information is available in MIMIC-Core.

Read more ...

MIMIC-IV-Hosp

The Hosp module provides all data acquired from the hospital wide electronic health record. This includes laboratory measurements, microbiology cultures, medication information, services provided, billed diagnoses and procedures, and so on.

Read more ...

MIMIC-IV-ICU

The ICU module contains information collected from the clinical information system used within the ICU. This includes highly granular information such as hour-to-hour vital signs, information about fluid management, and other charted observations.

Read more ...

MIMIC-IV-ED

The ED module contains data for emergency department patients including a triage assessment, nurse-validated vital signs, medicine reconciliation, and treatment information.

Read more ...

MIMIC-IV-CXR

MIMIC-cxr provides chest x-ray images and radiology reports for a subset of patients admitted to the emergency department.

Read more ...

MIMIC-IV-Note

The Note module contains deidentified free-text clinical notes for hospitalized patients.

Read more ...

MIMIC-III

MIMIC-III is an older version of MIMIC. It contains an older group of patients (ending in 2012), and a subset of the ICU and hospital information available in MIMIC-IV. We highly recommend researchers starting new studies to use the above modules in MIMIC-IV. We will maintain MIMIC-III for a short period to allow researchers to transition.

Read more ...

Questions?