AI In Electronic Medical Records

Open Access Teaching Case Developed for the Tech for Humanity Pathways Minor

Funded by the Andrew Mellon Foundation

Developed by Roan Parrish[1]

Background

Medical records are a core component of maintaining quality in medical care. However, recordkeeping has often been seen as a burden that takes up physicians' precious time. The shift from paper to electronic medical records has not solved the problem. Electronic health records (EHRs) are now ubiquitous, but they are associated with both improvements and problems in healthcare. Recent developments in generative AI have been applied to recordkeeping to mitigate some of these problems, particularly the huge amount of information presented to medical professionals and the significant time that practitioners and other workers spend entering data. However, these AI solutions introduce problems of their own, and may not address the underlying problems leading to physician burnout.

The 2009 HITECH Act introduced incentives for healthcare providers to adopt meaningful use of electronic medical records. In 2008, only 17% of office-based physicians and 9% of non-federal acute care hospitals were using EHRs. By 2021, that had surged to 78% of office-based physicians and 96% of hospitals having adopted a certified EHR.[2] However, the rapid adoption was far from seamless. The short timeframe meant the programs were generalized across hospitals, rather than localized and specific to their context. Procedures had to be adapted quickly and thus often inadequately. Change was implemented faster than policies directing the change. Common problems introduced by EHRs include alert fatigue (desensitization to warning

[1] March 1st, 2025

[2] Assistant Secretary for Technology Policy, "National Trends in Hospital and Physician Adoption of Electronic Health Records," Office of the National Coordinator for Health IT.

signs after encountering a massive number of them within a brief time), increased physician burn-out, and unclear or overwhelming information presentation.

Alert fatigue refers to how medical professionals become inured to the overwhelming amount of pings, alerts, and prompts presented every time they open the software. Important alerts are lost in the sea of reminders and notifications, some of which may not be clinical. Physicians can override alerts; but while this is often a clinical decision, overrides have more often become a way to dismiss common prompts.[3] Potential solutions have included stronger knowledge bases within the programs, such that alerts are less likely to be triggered and are more likely to be relevant. This may have the downside of requiring more detailed data entry, thus increasing the recordkeeping burden again. Another potential solution utilizes AI to interpret past information in the context of clinical guidelines, and predict whether an alert should be accepted or dismissed.[4] However, these models must be built on strong datasets, and the current trend of overworked health care practitioners dismissing annoying alerts must be considered in order to prevent building biases into the models.

Because of the push for all aspects of care to be captured in records, medical professionals spend more time updating their EHRs. However, the amount of time spent recordkeeping is significant enough to interfere with time spent with patients. Even worse, many physicians must spend unpaid time on documentation outside of clinic hours just to keep up, which decreases their overall well-being and contributes to burnout.[5] A common solution is to hire assistants to perform some tasks of routine documentation, particularly writing notes. Those transcribers can either be in the room with the physician and patient, or they can be remotely connected via an audio call. One troubling practice often used is to copy and paste chunks of text, which is likely to perpetuate inaccuracies and to overload the record with redundant notes that are more difficult to read later on.[6] Another solution is to have AI generate notes based on recorded audio and the data already entered into the EHR. The same generative AI can address the common problem of information bloat/overload in written patient records. Because patient notes are so long and thorough, it can be difficult for a physician to scroll through and skim quickly to find the

[3] Adam Wright et al., "Reduced Effectiveness of Interruptive Drug-Drug Interaction Alerts after Conversion to a Commercial Electronic Health Record," *Journal of General Internal Medicine* 33 (2018), 1868-1876.
[4] Siru Liu et al., "Leveraging Explainable Artificial Intelligence to Optimize Clinical Decision Support," *Journal of the American Medical Informatics Association* 31, no. 4 (2024), 968-974.
[5] A Jay Holmgren et al., "National Comparison of Ambulatory Physician Electronic Health Record Use Across Specialties," *Journal of General Internal Medicine* 39 (2024), 2868-2870.
[6] Robert Nagler Miller, "Copy, Paste, Repeat: Widespread EHR Practice Could Undermine Care," *American Medical Association*, June 29, 2017.

specific piece of information they need. Generative AI can produce summaries and interpret questions in real-time.

Generative AI still comes with the risk of “hallucinations,” or misinformation. This means that the created summaries may not be faithful to the original inputs. Incorrect summaries can seriously jeopardize patient health. Information presented in incorrect summaries may simply not be present in the original document; and the information may also *contradict* that of the original. The AI may also produce false logic, even using words that *are* present, by misinterpreting or misrepresenting content.[7] Many mistakes in generative AI summaries arise from ambiguous terms or unexpected terms in the notes. Clinical notes are filled with acronyms and alternative phrasings, which exacerbate the potential inaccuracies.[8] AI-produced summaries also risk taking data out of context. Physicians interpret data within the context of other factors, and AI summaries may not include all of the factors a physician would expect.

Case Study

The following is a fictional story.

Julie Snyder is concerned that her migraines have become more frequent and more intense, so she arranges an appointment with her primary care provider, Dr. Patricia Yates. After the front desk checks her in, a nurse performs an initial physical examination and talks Julie through her history and concerns. He sits at a computer next to Julie and enters her responses and her data into an electronic medical record. Shortly after the nurse leaves, Dr. Yates enters the office and takes the same seat in front of the computer. She glances through the freshly entered data, focuses on the complaint of migraines, and pulls up relevant records from previous visits. Because she doesn’t have much time to review the files, particularly with the patient waiting, Dr. Yates reads an AI-generated summary of the previous records about Julie’s migraines.

The physician turns to Julie and asks for permission to record their discussion for automatic generation of clinical notes. Julie hesitantly agrees, unsure about how this program works; her previous visits had a scribe in the corner to record notes instead. Julie and Dr. Yates spend the rest of the clinical encounter discussing the migraines and potential next steps, with the

[7] Prathiksha Rumale V et al., “Faithfulness Hallucination Detection in Healthcare AI,” *Proceedings of KDD 2024 Workshop - Artificial Intelligence and Data Science for Healthcare: Bridging Data-Centric AI and People-Centric Healthcare*, 2024.

[8] Savyasachi V. Shah, “Accuracy, Consistency, and Hallucination of Large Language Models When Analyzing Unstructured Clinical Notes in Electronic Medical Records,” *JAMA Network Open* 7, no. 8 (2024), e2425953.

physician only occasionally referring back to the computer. Julie leaves the visit with a new prescription for a daily prophylactic medication; when she stops by the pharmacy desk in the same building to pick up her medicine, Julie finds it ready to go because the record system automatically alerted the pharmacist technician. Julie quickly returns home.

Dr. Yates reviews the automatic transcription from the recording and checks the progress notes that the program's AI automatically generated. The transcription process is familiar, as the physician had been using a dictation service for many years, but the progress notes are a new function. Dr. Yates adds a few details that hadn't been picked up by the recording device, and then approves the notes before heading off to see her next patient. Now that the notes have been confirmed, the records program's AI generates medical codes about the diagnosis and treatment. These codes are then reviewed by a human medical coder to ensure accuracy and thoroughness. Codes are used for billing, for audits of clinical care, and for administrative planning.

A few days later, Julie is concerned she may be experiencing nausea as a side effect of the new medication. She decides to send her physician a message through the patient medical record portal. Later that day, Dr. Yates sits down to read her emails and messages from a number of patients. When she opens Julie's message, she can choose an AI-generated response to use as a template. The AI has already pulled data from Julie's record that it believes is relevant, including the most recent prescription. However, the AI has also pulled older data that is likely not applicable, so the physician deletes a portion of the response. Dr. Yates adds a paragraph explaining further details about the medication and how to decide if Julie should stop taking the medication. After reviewing the full response once more, the doctor hits "send", and turns to her next email.

Julie has a standing appointment every month at a specialist clinic. This clinic does not use AI in any part of their recordkeeping software. However, in order to stay updated on their patients, they do receive records from other clinics and hospitals, such as those created by Dr. Yates. The specialist clinic has no way of knowing if a record was generated by a human or an AI unless the provenance is indicated in the record itself, which is rare. Julie's appointment goes smoothly, and her specialist copies and pastes a standard progress note template into her record.

Processing Questions

1. What should the patient know in order to make an informed decision about an AI program recording their discussion?
2. How does the computer impact the perception of personal connection in the patient-physician relationship?
3. How is the presence of an in-person transcriber different from a remote scribe? How are scribes different from AI transcription and interpretation?
4. What might be the impacts of recording a patient in terms of what they are and are not willing to say? Does being recorded change what you are willing to say out loud?
5. If the AI record program is owned by a private company, what are their responsibilities in collecting and analyzing private health information?
6. What are the potential harms if the companies involved in this AI record are hacked? How are these harms different from existing potential harms?
7. Where are the major spots for error in the AI programs here?
8. What are the implications if an error occurs at each of these spots? Consider both minor and major consequences.

Thematic Reflection and Discussion

Privacy

Privacy is a primary concern with private health information. All clinical records, including any AI components, are regulated by HIPAA. The developer may require HIPAA authorization from each patient in order to include their data in the training set used to create the AI. Once the AI is in place, both the healthcare provider and the program company must maintain HIPAA compliance for all private health information. However, HIPAA's requirements may not be sufficient to ensure data security and privacy in rapidly evolving technologies.

1. What are the implications if a private company's healthcare AI is hacked?
2. Who owns the data produced through AI-mediated clinical care?

Transparency

AI algorithms are typically "black box", meaning the users and regulators have little insight into how algorithms interpret inputs that then produce the outputs used in healthcare. In this case, physicians don't know how the AI decides what should and should not be included in their summaries, and they do not have clear mechanisms for adjusting the AI's outputs to match their own needs. Hospital administration could run audits to see how often the AI produces material unfaithful to the original content, but they would not have real-time controls nor be able to tweak the program itself to address the results of the audits. Regulators are not given enough insight into black box algorithms to ensure safety and accountability.

1. How can developers make their AI records programs more responsive to hospital staff?
2. How can regulators design effective policies around black box algorithms?

Accountability and Responsibility

If a patient is harmed due to misinformation in an AI-generated summary, who would be held responsible? The physician is using a tool provided by a private company, but they are still responsible for their clinical decisions. Is the private company liable for their faulty product, or are AI hallucinations such a well-known phenomenon that the medical service providers should be considered as aware of the risks? What can be done to mitigate these risks? Regulations will probably be passed to determine accountability and responsibility, but these policies typically take years to be enacted.

1. How can patients navigate the medical system when they have no choices in whether or not AI is used in their care?
2. If a patient's medical records are transferred from an AI-using clinic to a non-AI-using clinic, how does accountability and responsibility transfer in the provision of care?

Connection

The medical provider-patient relationship is an integral component of healthcare. The provider may be a nurse, technician, or physician. Patients need to trust their provider. They need to feel heard. They need to know that their concerns are being taken seriously and that the provider is not brushing them off or dismissing their ailments. Physicians often complain that the computer gets in the way of that connection. The computer divides the physician's attention even as they are attempting to record what the patient is saying, or takes up time as the physician clicks madly, trying to find the right menu with the information they need right in the moment. The computer forms a literal, physical barrier between the physician and patient, or requires the physician to turn to the side rather than face the patient straight on. However, the computer also allows a patient to message their physician at any time, and to receive updates as soon as the physician has interpreted their lab results. Technology can both connect and distance.

1. How does an audio notes system change the computer's role in the clinical encounter? Consider both remote scribes and AI transcription. Would being recorded produce a chilling effect on patient speech?
2. How do AI-generated physician responses impact the physician-patient relationship?

References

Assistant Secretary for Technology Policy. "National Trends in Hospital and Physician Adoption of Electronic Health Records." Office of the National Coordinator for Health IT. https://www.healthit.gov/data/quickstats/national-trends-hospital-and-physician-adoption-electronic-health-records

Holmgren, A. Jay, Christine A. Sinsky, Lisa Rotenstein, and Nate C. Apathy. "National Comparison of Ambulatory Physician Electronic Health Record Use Across Specialties." *Journal of General Internal Medicine* 39 (2024): 2868-2870. https://link.springer.com/article/10.1007/s11606-024-08930-4

Khairat, Saif, Cameron Coleman, Paige Ottmar, Dipika Irene Jayachander, Thomas Bice, and Shannon S Carson. "Association of Electronic Health Record Use With Physician Fatigue and Efficiency." *JAMA Network Open* 3, no. 6 (2020): e207385. https://pmc.ncbi.nlm.nih.gov/articles/PMC7284310/

Liu, Siru et al. "Leveraging Explainable Artificial Intelligence to Optimize Clinical Decision Support." *Journal of the American Medical Informatics Association* 31, no. 4 (2024): 968-974. https://academic.oup.com/jamia/article/31/4/968/7610153

Liu, Siru, Allison B McCoy, Aileen P Wright, Scott D Nelson, Sean S Huang, Hasan B Ahmad, Sabrina E Carro, Jacob Franklin, James Brogan, and Adam Wright. "Why Do Users Override Alerts? Utilizing Large Language Model to Summarize Comments and Optimize Clinical Decision Support." *Journal of the American Medical Informatics Association* 31, no. 6 (2024): 1388–1396. https://academic.oup.com/jamia/article/31/6/1388/7624148

Nagler Miller, Robert. "Copy, Paste, Repeat: Widespread EHR Practice Could Undermine Care." *American Medical Association*, June 29, 2017. https://www.ama-assn.org/practice-management/digital/copy-paste-repeat-widespread-ehr-practice-could-undermine-care

Reddy Bongurala, Archana, Dhaval Save, Ankit Virmani, and Rahul Kashyap. "Transforming Health Care With Artificial Intelligence: Redefining Medical Documentation." *Mayo Clinic Proceedings: Digital Health* 2, no. 3 (2024): 342-347. https://www.mcpdigitalhealth.org/article/S2949-7612(24)00041-5/fulltext

Rumale V, Prathiksha et al. "Faithfulness Hallucination Detection in Healthcare AI." *Proceedings of KDD 2024 Workshop - Artificial Intelligence and Data Science for Healthcare: Bridging Data-Centric AI and People-Centric Healthcare*, 2024. https://openreview.net/pdf?id=6eMIzKFOpJ

Shah, Savyasachi V. “Accuracy, Consistency, and Hallucination of Large Language Models When Analyzing Unstructured Clinical Notes in Electronic Medical Records.” *JAMA Network Open* 7, no. 8 (2024): e2425953. https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2822301

Siwicki, Bill. “How Epic is Using AI to Change the Way EHRs Work.” *Healthcare IT News,* November 28, 2023. https://www.healthcareitnews.com/news/how-epic-using-ai-change-way-ehrs-work

Wright, Adam, Skye Aaron, Diane L. Seger, Lipika Samal, Gordon D. Schiff, and David W. Bates MD. “Reduced Effectiveness of Interruptive Drug-Drug Interaction Alerts after Conversion to a Commercial Electronic Health Record.” *Journal of General Internal Medicine* 33 (2018): 1868-1876. https://link.springer.com/article/10.1007/s11606-018-4415-9