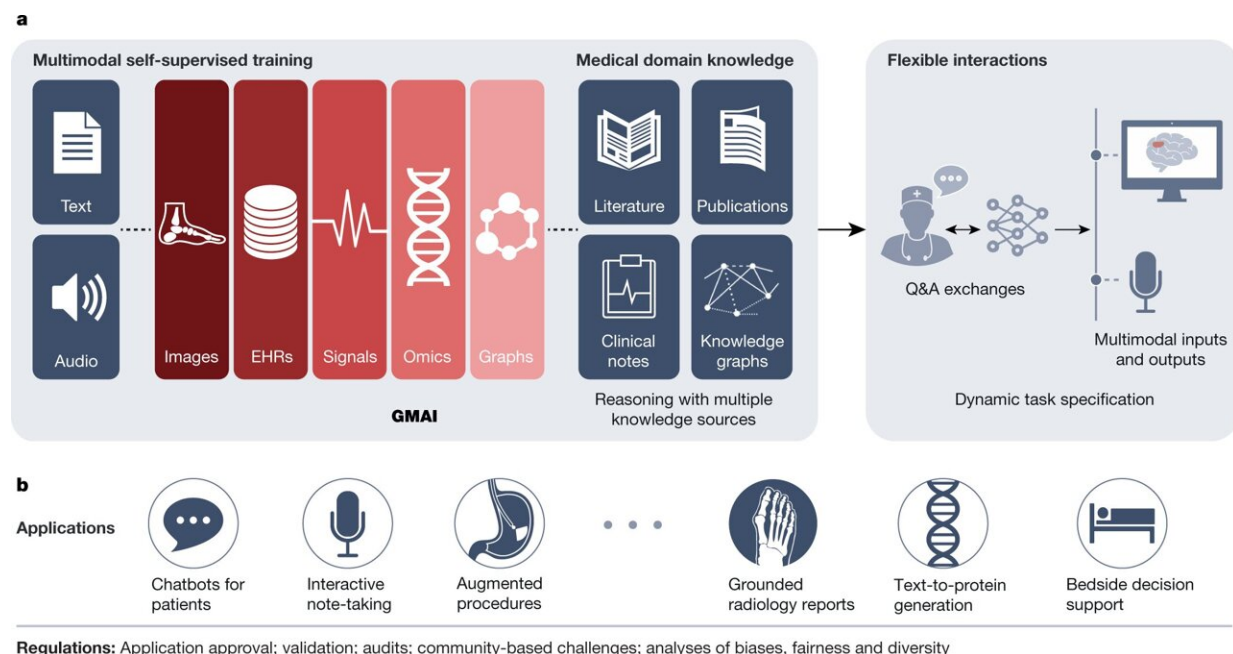


# The promise—and pitfalls—of medical AI headed our way

April 12 2023, by Corey Binns



Overview of a GMAI model pipeline. a, A GMAI model is trained on multiple medical data modalities, through techniques such as self-supervised learning. To enable flexible interactions, data modalities such as images or data from EHRs can be paired with language, either in the form of text or speech data. Next, the GMAI model needs to access various sources of medical knowledge to carry out medical reasoning tasks, unlocking a wealth of capabilities that can be used in downstream applications. The resulting GMAI model then carries out tasks that the user can specify in real time. For this, the GMAI model can retrieve contextual information from sources such as knowledge graphs or databases, leveraging formal medical knowledge to reason about previously unseen tasks. b, The GMAI model builds the foundation for numerous applications across clinical disciplines, each requiring careful validation and regulatory assessment.

Credit: *Nature* (2023). DOI: 10.1038/s41586-023-05881-4

A patient is lying on the operating table as the surgical team reaches an impasse. They can't find the intestinal rupture. A surgeon asks aloud, "Check whether we missed a view of any intestinal section in the visual feed of the last 15 minutes." An artificial intelligence medical assistant gets to work reviewing the patient's past scans and highlighting video streams of the procedure in real time. It alerts the team when they've skipped a step in the procedure and reads out relevant medical literature when surgeons encounter a rare anatomical phenomenon.

Doctors across all disciplines, with assistance from artificial intelligence, may soon have the ability to quickly consult a patient's entire medical file against the backdrop of all medical health care [data](#) and every published piece of medical literature online. This potential versatility in the doctor's office is only now possible due to the latest generation of AI models.

"We see a paradigm shift coming in the field of medical AI," said Jure Leskovec, professor of computer science at Stanford Engineering.

"Previously, medical AI models could only address very small, narrow pieces of the health care puzzle. Now we are entering a new era, where it's much more about larger pieces of the puzzle in this high stakes field."

Stanford researchers and their collaborators describe generalist medical artificial intelligence, or GMAI, as a new class of medical AI models that are knowledgeable, flexible, and reusable across many medical applications and data types. Their perspective on this advance is published in the April 12 issue of *Nature*.

Leskovec and his collaborators chronicle how GMAI will interpret varying combinations of data from imaging, electronic [health records](#), lab results, genomics, and medical text well beyond the abilities of concurrent models like ChatGPT. These GMAI models will provide spoken explanations, offer recommendations, draw sketches, and annotate images.

"A lot of inefficiencies and errors that happen in medicine today occur because of the hyper-specialization of human doctors and the slow and spotty flow of information," said co-first author Michael Moor, an MD and now postdoctoral scholar at Stanford Engineering. "The potential impact of generalist medical AI models could be profound because they wouldn't be just an expert in their own narrow area, but would have more abilities across specialties."

## Medicine without borders

Of the more than 500 AI models for clinical medicine approved by the FDA, most only perform one or two narrow tasks, such as scanning a chest X-ray for signs of pneumonia. But recent advances in foundation model research promise to solve more diverse and challenging tasks.

"The exciting and the groundbreaking part is that generalist medical AI models will be able to ingest different types of medical information—for example, imaging studies, lab results, and genomics data—to then perform tasks that we instruct them to do on the fly," said Leskovec.

"We expect to see a significant change in the way medical AI will operate," continued Moor. "Next, we will have devices that—rather than doing just a single task—can do maybe a thousand tasks, some of which were not even anticipated during model development."

The authors, who also include Oishi Banerjee and Pranav Rajpurkar from Harvard University, Harlan Krumholz from Yale, Zahra Shakeri

Hossein Abad from University of Toronto, and Eric Topol at the Scripps Research Translational Institute, outline how GMAI could tackle a variety of applications from chatbots with patients, to note-taking, all the way to bedside decision support for doctors.

In the radiology department, the authors propose, models could draft radiology reports that visually point out abnormalities, while taking the patient's history into account. Radiologists could improve their understanding of cases by chatting with GMAI models: "Can you highlight any new multiple sclerosis lesions that were not present in the previous image?"

In their paper, the scientists describe additional requirements and capabilities that are needed to develop GMAI into a trustworthy technology. They point out that the model must consume all of the personal medical data, as well as historical medical knowledge, and refer to it only when interacting with authorized users. It then needs to be able to hold a conversation with a patient, much like a triage nurse, or doctor to collect new evidence and data or suggest various treatment plans.

## **Concerns for future development**

In their research paper, the co-authors address the implications of a model capable of 1,000 medical assignments with the potential to learn even more. "We think the biggest problem for generalist models in medicine is verification. How do we know that the model is correct—and not just making things up?" Leskovec said.

They point to the flaws already being caught in the ChatGPT language model. Likewise, an AI-generated image of the Pope wearing a designer puffy coat is funny. "But if there's a high-stake scenario and the AI system decides about life and death, verification becomes really important," said Moor.

The authors continue that safeguarding [privacy](#) is also a necessity. "This is a huge problem because with models like ChatGPT and GPT-4, the online community has already identified ways to jailbreak the current safeguards in place," Moor said.

"Deciphering between the data and social biases also poses a grand challenge for GMAI," Leskovec added. GMAI models need the ability to focus on signals that are causal for a given disease and ignore spurious signals that only tend to correlate with the outcome. Assuming that model size is only going to get bigger, Moor points to early research that shows larger models tend to exhibit more [social biases](#) than smaller models. "It is the responsibility of the owners and developers of such models and vendors, especially if they're deploying them in hospitals, to really make sure that those biases are identified and addressed early on," said Moor.

"The current technology is very promising, but there's still a lot missing," Leskovec agreed. "The question is, can we identify current missing pieces, like verification of facts, understanding of biases, and explainability/justification of answers so that we give an agenda for the community on how to make progress to fully realize the profound potential of GMAI?"

**More information:** Michael Moor et al, Foundation models for generalist medical artificial intelligence, *Nature* (2023). [DOI: 10.1038/s41586-023-05881-4](#)

Provided by Stanford University

Citation: The promise—and pitfalls—of medical AI headed our way (2023, April 12) retrieved 6 May 2024 from <https://medicalxpress.com/news/2023-04-promiseand-pitfalls-of-medical-ai.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.