



Annual Review of Biomedical Data Science

Generative Artificial Intelligence: Implications for Biomedical and Health Professions Education

William Hersh

Division of Informatics, Clinical Epidemiology and Translational Science, Department of Medicine, School of Medicine, Oregon Health & Science University, Portland, Oregon, USA; email: hersh@ohsu.edu

Annu. Rev. Biomed. Data Sci. 2025. 8:355–80

The *Annual Review of Biomedical Data Science* is online at biomedata.sci.annualreviews.org

<https://doi.org/10.1146/annurev-biomedata-103123-094756>

Copyright © 2025 by the author(s). All rights reserved

Keywords

artificial intelligence, generative artificial intelligence, large language models, education, professional, educational assessment

Abstract

Generative artificial intelligence (AI) has had a profound impact on biomedicine and health, both in professional work and in education. Based on large language models (LLMs), generative AI has been found to perform as well as humans in simulated situations taking medical board exams, answering clinical questions, solving clinical cases, applying clinical reasoning, and summarizing information. Generative AI is also being used widely in education, performing well in academic courses and their assessments. This review summarizes the successes of LLMs and highlights some of their challenges in the context of education, most notably aspects that may undermine the acquisition of knowledge and skills for professional work. It then provides recommendations for best practices to overcome the shortcomings of LLM use in education. Although there are challenges for the use of generative AI in education, all students and faculty, in biomedicine and health and beyond, must have understanding and be competent in its use.



1. INTRODUCTION

The overall goal of this review is to present the key issues for generative artificial intelligence (AI) in the context of biomedical and health professions education. Biomedical and health professions students are defined broadly as any individuals who work professionally in healthcare, individual health, public health, and research. These include clinicians such as physician, nurses, pharmacists, etc., as well as those who do not provide clinical care, such as administrators, project managers, researchers, educators, and more. Also among these are those who work in biomedical and health informatics, data science, and related fields. Although the review is focused on AI in biomedicine and health, it draws on areas outside this broad domain where additional perspectives may be insightful.

AI has been defined in many ways, and this author prefers the notion that it consists of information systems and algorithms that are capable of performing tasks that are associated with human intelligence (1). AI is sometimes classified into two broad categories (2):

- Predictive AI uses data and algorithms to predict some output, e.g., diagnosis, treatment recommendation, prognosis, etc.
- Generative AI generates new output based on prompts, e.g., text, images, etc.

A large part of the modern success of AI is due to machine learning (ML), whose origins have been attributed to the Arthur Samuel in the 1950s and which has been defined as computer programs that learn without being explicitly programmed (3). The most success in applying ML to AI has come from so-called deep learning, which is based on many-layered neural networks (4). Predictive AI has typically used supervised ML, where models are trained with labeled or annotated data split into training and test data. On the other hand, generative AI systems are mostly trained with unsupervised ML due to the much larger sizes of training data and their generative, i.e., non-predictive uses, although some supervised ML takes place to avoid inappropriate generation, e.g., racist or sexist language (5).

Generative AI had been around for a number of years but came to the fore with the release of ChatGPT by OpenAI on November 30, 2022. Based on large language models (LLMs) trained using deep learning with massive amounts of digital text, generative AI systems have had a profound impact in biomedicine and health, including in education. A number of reviews describe LLMs generally (6, 7) as well as their applications in biomedicine (8–10).

This review focuses on generative AI, with a particular perspective on its implications for biomedical and health education. Some of the discussion may go beyond education, but only if there is some relevance to education. The review begins by describing the use of generative AI. It then discusses the successes and limitations of generative AI, especially those that apply to education. The review then covers issues specific to education, describing challenges and accomplishments, and then wraps up with future directions.

2. USAGE OF GENERATIVE ARTIFICIAL INTELLIGENCE

Since the introduction of ChatGPT and other systems, generative AI has been rapidly taken up by a large proportion of people around the world, including biomedical and health professionals, students, consumers, and others. A survey from mid-2023 found that 56% of college students stated they had used AI on assignments or exams, with 54% agreeing that the use of AI tools on college coursework counts could, in some instances, constitute cheating or plagiarism (11). A survey from mid-2024 of teachers, students, and parents found 49–52% of each group using generative AI frequently, 18–33% using it occasionally, and less than a quarter of each group

Table 1 Use of sources for health information and advice

| Source | Every day | At least once a week | At least once a month | Occasionally | Never |
|----------------------------------|-----------|----------------------|-----------------------|--------------|-------|
| Social media | 46% | 10% | 4% | 18% | 22% |
| Internet search | 42% | 21% | 9% | 21% | 7% |
| Artificial intelligence chatbots | 6% | 7% | 4% | 19% | 63% |

Data from Reference 15.

stating they had never used it (12). Both of these surveys noted positive views of AI but also a lack of consistent or comprehensive policies by educational institutions and faculty within them.

Other surveys have looked at larger populations and noted similar large-scale uptake of generative AI. One analysis from August 2024 found that 39% of working-age adults use generative AI (13). It also found that more than 24% of such individuals used it in their work at least once in the week, and approximately 10% used it every day at work. The authors of the survey noted that the adoption of generative AI has been faster than even adoption of personal computers or the Internet. Another survey, this one of over 1,000 physicians from the United Kingdom, found that 20% reported using generative AI tools in clinical practice (14). Among those who used such tools, 29% reported using it to generate documentation after patient appointments and 28% to develop a differential diagnosis.

A survey of 2,428 US adults of varying age, geographic location, and race/ethnicity carried out by the Kaiser Family Foundation found that approximately two-thirds of respondents reported some use or interaction with AI (15). However, use of AI chatbots for health information and advice was much lower than for social media or Internet search (**Table 1**).

3. RESULTS OF GENERATIVE ARTIFICIAL INTELLIGENCE IN BIOMEDICINE AND HEALTH

An avalanche of research has been published on the results of generative AI applied to biomedical and health tasks. Many of these studies have focused on answering questions, especially those on examinations, including medical board tests. Others have expanded to solving clinical cases or carrying out other clinical tasks, although few involve actual users and almost none have been implemented in real-world settings. Nonetheless, the activities considered in these studies impact education in that they are related to tasks used in educational activities, including student practice and assessments. The following categories of research are described in this section: medical board examinations, other medical examinations, answering clinical questions, solving clinical cases, assessing clinical reasoning, summarization tasks, predictive tasks, performance in graduation courses in biomedicine and health, and performance in academic courses beyond biomedicine and health.

3.1. Medical Board Examinations

One of the most widely used datasets for LLM evaluation comes from a set of sample questions from the US Medical Licensing Exam (USMLE) (16). This multiple-choice question (MCQ) exam is taken in three steps during the second and fourth years of medical school and then in the first year out of medical school. The dataset, called MedQA, includes 12,723 questions in English (and others in Chinese). An early and highly publicized success of the original ChatGPT using the GPT-3.5 LLM on the MedQA dataset was to achieve a score of over 60%, which would equate to a passing grade (17). This level was surpassed by other LLMs, including GPT-4 (18) and Google's Med-Gemini (19). The most recent leader in the USMLE arms race is OpenAI's o1 model, which

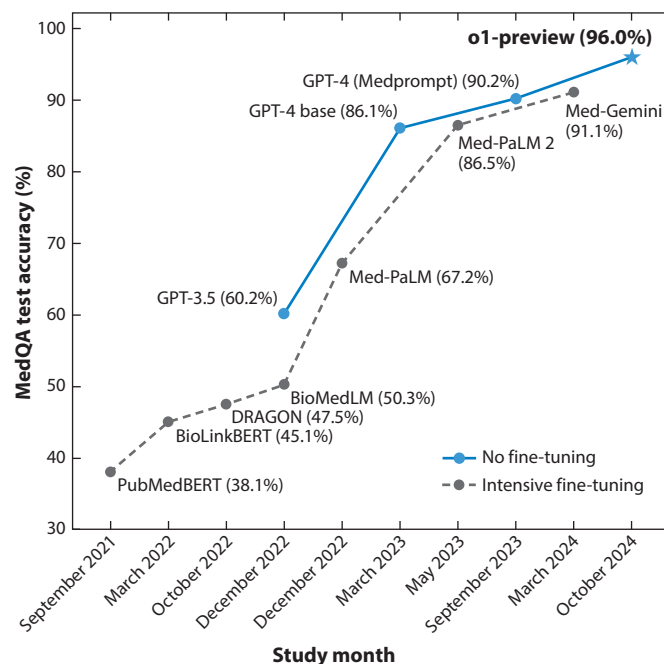


Figure 1

Incremental progress of large language models on the US Medical Licensing Exam's MedQA dataset over time (20, 21). Figure adapted with permission from Reference 20.

scored 96.0% (20, 21) (**Figure 1**). LLMs have even been found to perform well on portions of the USMLE exam covering soft skills, e.g., communication skills, ethics, empathy, and professionalism (22).

This has led to the assessment of LLMs on other US-based medical board exams, with similar success being found, among others, in the fields of radiology (23, 24), neurosurgery (25), and clinical informatics (26). LLM success on board exams is not limited to the United States. In Israel, the scores of all resident physicians who completed board exams in the specialties of pediatrics, internal medicine, psychiatry, obstetrics/gynecology, and general surgery in 2022 were compared to those of GPT-3.5 and GPT-4 (27). As seen in **Figure 2**, GPT-4 passed the exams in all of the specialties, scoring comparably in internal medicine and general surgery, better in psychiatry, and inferiorly in pediatrics and obstetrics/gynecology.

3.2. Other Medical Examinations

Success with generative AI has been shown with other kinds of physician examinations. On an objective structured clinical examination, which provides a series of questions simulating a clinical encounter, ChatGPT-3.5 scored better than a group of physicians in Singapore in a simulation of assessment for membership in the Royal College of Obstetricians and Gynecologists (28). Likewise, ChatGPT-4 achieved 81% correct on the Ophthalmology Knowledge Assessment Program (29) and 73.3% on the Nephrology Self-Assessment Program (30).

Many medical specialties give in-training exams to residents and fellows in training. On a family medicine in-training exam, ChatGPT-4 scored 86.5% and outperformed third-year resident national averages in one study (31) and scored 84% on another (32). The latter study also found that




Figure 2

LLM success versus physicians completing residency training in five specialties in Israel. The range is shown for all test takers, human and LLM. The dotted line is the passing level for the exam for each specialty. Abbreviation: LLM, large language model. Figure adapted with permission from Reference 27.

GPT-4 was able to integrate new information and carry out self-correction when needed. Likewise, on a radiology in-training exam, GPT-4 achieved results comparable to between second-year and third-year radiology residents (33). On this exam, performance on image-based questions was significantly poorer at 45.4% compared to text-only questions at 80.0%. As has been seen in some other studies, when questions were given again, GPT-4 chose a different answer approximately 25% of the time. Fine-tuning the LLM did not improve accuracy in this study.

3.3. Answering Clinical Questions

The success of generative AI is not limited to exams. Other research has assessed its ability to answer clinical questions in a variety of medical disciplines. One of the most studied applications areas has been cancer. One study assessed ChatGPT output for concordance with National Comprehensive Cancer Network treatment guidelines for breast, prostate, and lung cancer. The study found an overall concordance 61.9% of the time, with 34.3% of outputs recommending one or more nonconcordant treatments. The study also found that responses were hallucinated, i.e., were not part of any recommended treatment, in 12.5% of outputs (34). Another study assessed several different LLMs and found varying performance on different categories of oncology questions, with ChatGPT-4 the only LLM to score higher than 50% (35).

An additional study assessed the ability of GPT-4 to answer clinically relevant questions regarding the management of patients with pancreatic cancer, metastatic colorectal cancer, and hepatocellular carcinoma from clinical practice guidelines (36). Results found that the addition of information for the use of retrieval-augmented generation (RAG), a method that allows new information to be added to already trained LLMs, resulted in correct responses in 84% of cases compared to only 57% without the use of RAG.

The success of LLMs is not limited to cancer. One study of 284 questions developed by physicians found that ChatGPT-4 had highly accurate and complete answers and performed better than ChatGPT-3.5 (37). Another study looked at ChatGPT-3.5 answering MCQs about human genetics (38). ChatGPT-3.5 responses were found to be correct comparably to human answerers (68.2% versus 66.6%, respectively). It was also found that both ChatGPT and humans performed better on memorization-type questions than on critical thinking questions. Another interesting finding was that, when asked the same question multiple times, ChatGPT provided

different answers 16% of the time, including for both initially correct and incorrect answers, and gave plausible explanations for both correct and incorrect answers.

An additional study assessed GPT-4 for concordance with recommendations made by 12 physicians from a hospital consultation service (39). There were 31 of 66 questions that had a majority (i.e., more than 6) of physicians rating concordance. Of these, the responses to 13 questions were designated concordant, 15 were discordant, and 3 were unable to be assessed. The responses from GPT-4 were found to largely be devoid of overt harm, but fewer than 20% of the responses agreed with the answer from the consultation service. It was noted that some responses contained hallucinated references.

Other studies have focused on specific types of questions for clinicians. One looked at the ability of GPT-4 to appropriately perform probabilistic reasoning by comparing its performance with a large survey of human clinicians (40). The LLM was more accurate than human clinicians in determining pretest and posttest probability after a negative test result for five cases but did not perform as well after positive test results. A couple of additional studies have focused on answering clinical questions using general LLMs further tuned with clinical information resources. One system called Almanac, an LLM framework using RAG from curated medical resources for medical guideline and treatment recommendations, was found to obtain significant improvement in performance compared with standard LLMs on measures of factuality, completeness, user preference, and adversarial safety (41). Another system based on RAG and called ChatRWD outperformed the commercial system OpenEvidence as well as several general LLMs on questions seeking additional clinical evidence (42).

There are many other medical question-answering datasets that have been developed and used to compare generative AI systems. One recent analysis used a large number of them to compare the new OpenAI o1 LLM that makes use of a new type of prompting called chain of thought prompting, where prompts to the LLM describe the task iteratively, and found that it surpassed GPT-4 in accuracy by an average of 6.2% and 6.6%, respectively, across 19 datasets and two newly created complex question-answering scenarios (43). However, the study also noted several weaknesses of the LLMs, including hallucination, inconsistent multilingual ability, and discrepant metrics for evaluation. We can see from all of these studies that generative AI is very good at answering questions on medical examinations, yet it is far from perfect and does suffer from incorrect and/or inconsistent answers.

3.4. Solving Clinical Cases

Some analyses of generative AI have gone beyond answering questions to assessing the ability to solve clinical cases. One line of research has involved a collection of clinical vignettes that were originally developed in the mid-2010s to assess the performance of clinical symptom checkers (44). Most systems using them performed poorly at the time, and certainly far worse than physicians, who averaged approximately 72% accuracy. One study assessed ChatGPT-3.5 with this dataset for first-pass diagnostic and triage decision accuracy, finding that ChatGPT-3.5 identified illnesses with 75.6% first-pass diagnostic accuracy and 57.8% triage accuracy (45). This study also found that ChatGPT was useful for generating new vignettes written for those with both high and low health literacy levels. Another study found that ChatGPT-3.5 identified the correct diagnosis in the top three diagnoses for 88% of cases, compared to 54% for lay individuals and 96% for physicians (46). ChatGPT-3.5 was also found to triage 71% of cases correctly, similar to lay individuals, with both worse than physicians, who triaged correctly for 91% of the cases.

Other studies have used the well-known clinicopathologic conferences featured in the *New England Journal of Medicine* (NEJM). These cases are deemed to be challenging and often stump even Harvard Medical School faculty. One study found that GPT-4 provided the correct diagnosis

within the differential diagnosis in 64% of the 70 cases assessed and as the top diagnosis in 39% of cases (47). Another study found GPT-4 to be correct for 57% of 38 cases assessed, which was better than almost all online readers who answered the challenge (48). A follow-on analysis compared performance for GPT-4 with cases that were newer and older than the September 2021 training date of GPT-4, aiming to control for possible data leakage, i.e., the early cases potentially being in the training data. This analysis found comparable results in newer and older cases.

Another clinical problem-solving study assessed Google's Med-PaLM 2, which had been optimized for diagnostic clinical reasoning (49). Unlike earlier LLM studies, this one employed real physicians solving 302 NEJM clinicopathologic conferences. Generalist physicians were given versions of cases that were redacted for diagnostic testing and the final diagnosis. The physicians were asked to generate a differential diagnosis when randomized to one of two conditions, one having access to standard online search tools or being provided output from Med-PaLM 2. Specialist physicians were provided access to the gold standard–evaluated differential diagnosis lists and asked to evaluate the lists for inclusion of the final diagnosis, comprehensiveness of the differential diagnosis, and appropriateness of the differential diagnosis. The overall best diagnosis in the top 10 of the differential diagnoses came from Med-PaLM 2 only (59.1%), followed by generalist physicians with Med-PaLM 2 (51.7%), generalist physicians with access to conventional search systems (44.4%), and unassisted generalist physicians (33.6%). Also of note, this Google LLM was found to exceed the performance of GPT-4.

An additional dataset used for clinical problem-solving has been a collection of 36 clinical vignettes from the MSD Clinical Manual, formerly known as the Merck Manual (50). In this study, ChatGPT was found to achieve overall correctness on all questions for all cases 71.7% of the time. It performed best for answering the final diagnosis and lowest for generating an initial differential diagnosis. The overall accuracy was lower for diagnostic and management questions than for diagnosis questions. There was no variation in answer correctness by the age, gender, or acuity of the patient in the vignette.

Some studies have looked at solving clinical cases in specific settings. One study from a Dutch emergency department (ED) retrospectively reviewed notes and entered them into ChatGPT-3.5 or ChatGPT-4 (51). For generating differential and leading diagnoses, ChatGPT-4 performed comparably to physicians. Another notable finding was that submitting the identical query to ChatGPT-3.5 or ChatGPT-4 three different times had the same leading diagnosis only 60% of the time and overlap of all of the differential diagnoses only 70% of the time.

ChatGPT-4 output has also been found to align well with accepted guidelines for managing mild and severe depression, without showing the gender or socioeconomic biases sometimes observed among primary care physicians (52). ChatGPT-4 has also been found to provide highly accurate (98.9%) medication recommendations as a second opinion in a dermatology treatment setting, but its reliability and comprehensiveness were deemed by researchers to need refinement for greater accuracy (53).

Another study prompted ChatGPT-4 with symptoms of 194 diseases in the Mayo Clinic Symptom Checker, and ChatGPT-4 achieved 78.8% accuracy in making the correct diagnosis (54). The performance varied by clinical specialty, with best results for dentistry, endocrinology, and infectious diseases and worst results for pulmonology, dermatology, and emergency medicine.

3.5. Assessing Clinical Reasoning

A number of studies have been based on instruments that assess clinical reasoning. In one study of clinical reasoning using 20 patient cases, GPT-4 was found to perform comparably to attending physicians and residents in diagnostic accuracy, correct clinical reasoning, and cannot-miss diagnosis inclusion (55). In another pair of studies of diagnostic reasoning, physicians were assessed on



a set of clinical vignettes developed to assess clinical decision support systems in the 1990s (56). Physicians were randomized for different vignettes to have conventional information resources with or without the addition of GPT-4 and assessed using an instrument to assess diagnostic reasoning (57). Physicians using GPT-4 scored comparably to those not using it (76% versus 73%), although those using ChatGPT-4 showed a tendency to solving cases faster (565 versus 519 seconds). GPT-4 alone scored much higher at 92%. In a similar randomized vignette study assessing clinical management decisions, physicians scored 6.5% higher using an LLM compared to those using conventional resources, but in this instance, GPT-4 alone did not do better (58).

Another study of clinical reasoning was based on a framework developed to simulate a realistic clinical setting based on diagnostic accuracy, adherence to diagnostic and treatment guidelines, consistency in following instructions, ability to interpret laboratory results, and robustness to changes in instruction, information quantity, and information order (59). Using the Medical Information Mart for Intensive Care (MIMIC) database of 2,400 real patient cases and four common abdominal pathologies, it was found that LLMs were not able to accurately diagnose patients across a number of different pathologies, performing significantly worse than physicians. The LLMs did not follow diagnostic or treatment guidelines and were not able to correctly interpret laboratory results. The authors also concluded that these LLMs would not be able to easily integrate into existing workflows because they frequently did not follow instructions and were sensitive to both the quantity and order of information. One limitation of this study was that most of the known best-performing LLMs were not able to be used due to licensing restrictions that prohibited MIMIC data from being used in commercial versions of LLMs.

A Google LLM that has been studied with real users is the Articulate Medical Intelligence Explorer (AMIE). An initial study applied AMIE in a randomized controlled trial that compared primary care physicians and the AMIE output, judged by patient actors and specialist physicians (60). The study was limited by the use of text-based dialogue for human–system interaction, but AMIE was found to outperform primary care physicians in history taking, diagnostic accuracy, management reasoning, communication skills, and empathy. A second study developed a real-world dataset of 204 cardiology cases that included reports from electrocardiograms, echocardiograms, cardiac magnetic resonance imaging, genetic tests, and cardiopulmonary stress tests (61). A 10-domain evaluation rubric was developed and used by cardiologists to evaluate the quality of diagnosis and clinical management plans for the cases entered into AMIE. The recommendations from AMIE were rated superior to general cardiologists for 5 of the 10 domains and equivalent for the others.

Another study of clinical problem-solving was a blinded observational comparative study conducted in a primary care setting in Sweden (62). Responses from GPT-4 and real physicians to cases from a family medicine specialist examination were scored by blinded reviewers. In these complex primary care cases, GPT-4 and the follow-on GPT-4o performed worse than average human physicians. Recognizing the need to assess clinical reasoning with more complex dialogues and less with MCQs, the Conversational Reasoning Assessment Framework for Testing in Medicine (CRAFT-MD) dataset was developed (63). This dataset focuses on natural dialogues, using simulated agents to interact with LLMs in a controlled environment. The CRAFT-MD framework showed that LLMs performed notably worse in conversational settings compared to in examination-based evaluations. The authors noted that more realistic testing approaches must be developed before LLMs can be safely integrated into clinical workflows and proposed a set of recommendations to align LLM evaluations with real-world clinical practice.

Although all of the studies presented so far have been based on textual data, some studies have assessed the use of foundational LLM models that include images and clinical tasks using them. In one study of image diagnosis using cases from the JAMA Clinical Challenge and NEJM Image

Challenge databases, OpenAI's GPT-4V demonstrated better image-interpretation accuracy than unimodal LLMs (64). Another study based on NEJM Image Challenges found that GPT-4V performed comparably to human physicians regarding multiple-choice accuracy (81.6% versus 77.8%) (65). GPT-4V also performed well in cases where physicians answered incorrectly, with approximately 78% accuracy. However, GPT-4V was also found to provide flawed rationales in cases where it made the correct final choices (35.5%), most prominently in image comprehension (27.2%).

An additional study combined PaLM with radiology reports along with an image encoder and enabled the detection of five findings on chest X-rays: atelectasis, cardiomegaly, consolidation, pleural effusion, and pulmonary edema (66). Another study used a model for images and text that outperformed image-only and non-unified models for prediction of adverse events in pulmonary disease (67). An additional study combined a locally developed LLM and an image classifier trained on prior chest X-rays paired with reports (68). The resulting system was assessed with new chest X-rays in an ED and found to produce draft reports with clinical accuracy and textual quality that were comparable to on-site radiologist reports and provided higher textual quality than reports from off-site teleradiologists. Finally, another study used the LLaVA-Med LLM with image-report pairing to correctly answer questions from a visual question-answering dataset better than previous systems (69).

One role for generative AI in radiology may be to assist radiologists in their reports. One study noted the ability of GPT-4 to identify missed diagnoses in preliminary reports of radiology trainees (70). Another study found that GPT-4 detected errors in radiology reports (e.g., omissions, insertions, spelling, side confusion, and more) comparably to radiologists and did so much faster than them (average of 3.5 versus 25 seconds) (71).

3.6. Summarization Tasks

As education often involves summarization of scientific articles, medical records, and other documents, LLMs have been assessed in summarization tasks. One study assessed GPT-4 at providing feedback to authors of computational biology scientific papers (72). Assessing the PDFs of papers, the LLM feedback was found to have overlap comparable to that of humans and was deemed higher for poorer-quality papers. Over half of the papers' authors found the generated feedback to be helpful or very helpful, and 82% found it more beneficial than feedback from at least some human reviewers. Another study looked at summaries of 140 evidence-based journal abstracts generated by ChatGPT-3.5; the summaries were 70% shorter than the mean abstract length and were found to have high quality, high accuracy, and low bias (73). An additional system for synthesis of knowledge used a combination of gathering evidence, traversing citations, and synthesizing the results to perform better than humans in question-answering, summarization, and contradiction detection (74).

ChatGPT-4 has also been assessed in the generation of lay summaries of scientific abstracts for a national clinical study recruitment platform, ResearchMatch (75). ChatGPT-4 achieved 95.9% accuracy and 96.2% relevance across 192 summary sentences from 33 abstracts of clinical studies. A total of 85.3% of 34 volunteers rated ChatGPT-generated summaries as more accessible and 73.5% rated them more transparent than the original abstract. None of the summaries were rated as harmful by clinical experts.

3.7. Predictive Tasks

While this review distinguished at the onset a difference between predictive and generative AI, the latter has been found in a number of instances to have predictive success as well. GPT-4 was

shown to achieve accuracy of predicting cardiovascular disease comparable to the well-known Framingham predictive model with data from the UK Biobank and the Korean Genome and Epidemiology Study (76). Other studies have assessed GPT-4 for predictive abilities in the ED, finding suboptimal ability to predict acuity of patients in the ED (77) or predict admissions from ED to the hospital (78). However, better success has been found in identifying missed diagnoses in radiology residents' reports (70) and in increasing the accuracy of rare disease diagnosis (79).

3.8. Performance in Graduate Courses in Biomedicine and Health

Although the majority of biomedical and health applications of LLMs discussed so far have focused on medicine, some applications have focused on graduate-level study. One analysis compared LLM and student performance on graduate-level examinations in biomedical sciences (80). The performance of GPT-4 on nine examinations from courses in topics such as genomics, microbiology, and cellular and molecular biology was compared with actual student performance. GPT-4 was found to exceed the student average on seven of nine exams and for all student scores for four exams. GPT-4 performed very well on fill-in-the-blank, short-answer, and essay questions as well as questions on figures sourced from published manuscripts. It performed less well on questions with figures containing simulated data or those requiring hand-drawn answers. The analysis found that some model responses included hallucinations.

Another analysis compared student knowledge-assessment scores in a large online introductory course in biomedical and health informatics by prompting six LLM systems as they would be used by typical students (81). The six LLM systems were prompted to answer 10 MCQs each from the 10 units of the course and 33 final exam questions in this course that is taken by graduate, continuing education, and medical students. Scores for 139 students (30 graduate students, 85 continuing education students, and 24 medical students) who took the course in 2023 were compared to those of the LLM systems. All of the LLMs did well enough to achieve a passing grade in the course and scored between the 50th and 75th percentiles of students, with Google's Gemini scoring highest across all assessments (see **Figure 3**). The performance of the LLMs raised challenges




Figure 3

Student scores for 25th, 50th (median), and 75th quartile of performance (*thinner blue, purple, and green lines, respectively*) versus the best-performing large language model, Gemini Pro (*thicker gray line*), in the individual and aggregate unit assessments and the final examination for a large, online introductory biomedical and health informatics course. Abbreviation: MCQ, multiple-choice question. Figure adapted from Reference 81.

about student assessment in higher education, especially in courses that are knowledge-based and online.

Another study from health informatics looked at the use of GitHub CoPilot in a programming course (82). GitHub is a computer code repository system that has been integrated with an OpenAI LLM that aims to assist in the writing of computer code. The system was evaluated in a health informatics programming course for two types of programming problems, one for database queries in structured query language (SQL) and the other for computational tasks using the Python programming language. In general, the generated solutions worked well for simple and straightforward SQL and Python tasks but less well for more complex ones. It was also noted that some solutions were correct but did not take the most efficient approach to programming.

3.9. Performance in Academic Courses Beyond Biomedicine and Health

Generative AI systems have also been assessed in education beyond biomedicine and health. One report describing the release of a new open-source LLM, Llama 3, noted that, like other major LLMs, it was capable of passing high school advanced placement and college prep tests (83). At the college undergraduate level, examinations for five courses in psychology at a university in the United Kingdom were written by GPT-4 and achieved grades higher than the average of students taking the exams in the courses (84). It was also found that 94% of the AI submissions went undetected. Success has been achieved at even more advanced academic levels, such as on the so-called Graduate-Level Google-Proof Q&A Benchmark, a challenging set of questions from biology, chemistry, and physics (85). The OpenAI o1 LLM was found to exceed the scores of PhD-level students on this dataset (86).

One area where generative AI has had major impact has been computer science. In an overview, it was noted that LLMs are very capable computer programmers, raising questions about the teaching and even the practice of computer programming going forward (87). This author noted that LLMs are capable of generating solutions to problems typical of introductory programming courses, raising concerns around potential student overreliance and misuse. A comprehensive analysis of many different Python programming tasks found that GPT-4 could create computer code even for complex tasks but still required humans to ensure the validity and accuracy of the code (88). Another author noted that computer science educators at all types of institutions must monitor changes and adjust curriculum accordingly (89).

The success of LLMs in academic tasks has been demonstrated in other disciplines. For law students, GPT-4 was found to modestly improve students' performance in a number of legal tasks and to increase the speed of completion markedly (90). Likewise, LLMs have been found to be able to create functional pipelines in data science tasks (91, 92). Another study found that GPT-4 was able to solve novel and difficult tasks that span mathematics, coding, vision, medicine, law, psychology, and more, without the need for any special prompting (93).

Some studies have looked at the use of LLMs outside the academic arena, but their mixed results have implications for education. One study looked at a creative divergent thinking task called the Alternative Uses Task and found that ChatGPT-3.5 and 4 outperformed average but not the best humans (94). Another study found that ChatGPT-4 exceeded other LLMs at a variety of general human language tasks but performed less well on reasoning tasks and was still prone to hallucinations (95). An additional study that assigned writing tasks to 453 college-educated professionals found a 40% decrease in time and an 18% improvement in quality for the half of subjects assigned to use ChatGPT (96). In business, GPT-4 has been found to outperform analysts in the ability to predict earnings changes in companies (97, 98). In software engineering, three

randomized controlled trials showed an average of 26% productivity gains, with the gains higher for workers with less experience (99).

Another study at a global management consulting firm randomized consultants to using ChatGPT-4 or not in their work (100). Those assigned to ChatGPT-4 were found to be more productive, completing an average of 12.2% more tasks, and more quickly completing tasks 25.1% of the time. Consultants assigned ChatGPT-4 were also found to produce higher-quality results, measured to be more than 40% higher quality compared to control group. For some tasks, however, such as combining qualitative and quantitative data, ChatGPT-4 performed less well, leading the authors to note there was a so-called jagged technological frontier, where some tasks were easily done by AI but others not.

4. LIMITATIONS OF GENERATIVE ARTIFICIAL INTELLIGENCE

The limitations of generative AI are well-known and are reviewed here from the perspective of education. From the earliest days of ChatGPT, it was seen that LLMs sometimes hallucinate. Indeed, the Dictionary.com 2023 Word of the Year was hallucinate, mostly based on its connection to ChatGPT and other generative AI (101). This section describes the following limitations of generative AI: factuality of LLMs; citations and attribution; data bias, leakage, and drift; detecting use of generative AI; and real-world and safe use.

4.1. Factuality of Large Language Models

Going beyond hallucinations and confabulations, Augenstein has described “factuality challenges” (102) for LLMs, which include the following:

- Undersourcing, i.e., lacking credible references to statements made
- Truthfulness, i.e., their hallucinations
- Speaking with a confident tone and fluent style, which may mislead users who are less familiar with the facts about a given topic
- Allowing direct use by users; ease of access
- Halo effect from being knowledgeable in domains other than the one of current interest to a user
- Public perception that LLMs are all-knowing
- Unreliable evaluation, leading to excess positivity about their functionality

4.2. Citations and Attribution

A concern related to factuality is the tendency for LLMs to generate fabrications and errors in citations, an important aspect of scientific and scholarly communication. One study prompted ChatGPT to produce short literature reviews on 42 multidisciplinary topics. Analysis found that 55% of GPT-3.5 citations and 18% of GPT-4 citations were fabricated. In addition, 43% of real GPT-3.5 citations and 24% of real GPT-4 citations included substantive errors (103).

A related concern is that current LLMs are not necessarily good at attributing their information with references. One study looked at reference validity and attribution for several LLMs (104). It was found that GPT-4, as used in Microsoft Copilot, had very high validity of URL sources cited, with other LLMs performing more poorly. But even GPT-4 in Copilot only provided statement-level support for 70% of its assertions and response-level support for 54%, with the others faring worse. Even Copilot failed to cite any sources for approximately 20% of prompts, with others having higher rates of non-citation. Another issue that this study raised was attribution sources behind paywalls that users might not be able to access.

These results led one author to note that, from the perspective of biomedicine and health, especially the academic use of searching, there may be concerns regarding authoritativeness, timeliness, and contextualization of search results. In other words, we often search not just to find an answer, but also to find out where the answer came from, who wrote it, and what methods they used (105).

4.3. Data Bias, Leakage, and Drift

LLMs have also been documented to perpetuate bias that is present in many information sources that are used for training LLMs, including scientific literature and other information on the Internet. One study assessed bias by asking eight clinical questions of four different LLMs, such as estimated glomerular filtration rate, lung capacity, and pain threshold (106). The four LLMs—ChatGPT-3.5, Bard, Claude, and GPT-4—all were found to recapitulate what the authors called harmful, race-based medicine.

Another study analyzed standardized clinical vignettes from the publication NEJM Healer (107). GPT-4 was found to be more likely to include diagnoses that stereotyped certain races, ethnicities, and genders, and did not model appropriate demographic diversity of various medical conditions in the vignettes. Outside of medicine, one study assessed AI-generated content (AIGC) produced based on headlines in the *New York Times* and *Reuters* and found substantial gender and racial biases (108). AIGC generated by each LLM exhibited notable discrimination against female and Black individuals. Among the LLMs, AIGC generated by ChatGPT demonstrated the lowest level of bias, and ChatGPT was the only LLM able to decline content generation when biased prompts were entered.

Additional concerns for LLMs include issues around the data used to train them. One issue is data leakage, which results when there is contamination of model input features with outcome information. Or put another way, when training data make it into the test data and potentially influence the result, this may artificially improve the metric being used to evaluate the model. One analysis found at least 294 papers affected by leakage across 17 different scientific disciplines and identified eight reasons for data leakage. Whatever the reason, the usual cause was training data making their way into or somehow otherwise influencing test data, potentially biasing the performance results from the model (109).

A related concern is dataset shifts due to changes that may occur in technology, population and setting, and clinician or patient behavior (110). Diseases may change over time as well, with a good example provided by COVID-19, which is a different disease now than when the pandemic began in 2020 (111). Another type of information that changes over time is genetic information, which may be due to new associations found in research or new technologies used. One study developed a chained, two-prompt GPT-4 sequence from a training set of 45 article–variant pairs for the automated classification of functional genetic evidence in the scientific literature. The first prompt asked GPT-4 to supply functional evidence in a given article for the variant of interest or indicate the lack of such evidence. For articles in which GPT-4 designated functional evidence, a second prompt asked it to classify the evidence into categories of pathogenic, benign, or intermediate/inconclusive. A test set of 72 manually classified article–variant pairs found substantial variability over 2.5 months in the results both within prompts entered in rapid succession and across days (112).

4.4. Detecting Use of Generative Artificial Intelligence

One desired feature of LLMs in many applications, including those in education, is the ability for people or computer algorithms to detect their use. Certainly in education, it would be valuable to be able to discern the use of LLMs, especially in situations where educators might not want them

to be used. Unfortunately, research shows that both machine and human detection of LLM usage is inconsistent at best (113).

One system that focused on research abstracts from scientific journals was shown to have a 98% rate of detection accuracy and perform much better than humans (114). Another pair of studies found that an ML model could distinguish scientific writing from ChatGPT output (115), including its focused use in chemistry journals (116). However, another study found that light paraphrasing undermined generative AI detectors (117). A further evaluation of 11 Web-based detectors found that simple modifications, such as introduction of minor grammatical errors and substitution of Latin with similar Cyrillic letters, undermined detectors (118). Another study found that the already-low accuracy rates (39.5%) of LLM detectors showed further reductions in accuracy (17.4%) when faced with manipulated content, with some techniques proving more effective than others in evading detection (119). Another concern for LLM detectors is that some systems are more likely to classify non-native English writing as AI generated (120).

Humans are not good detectors of LLM writing either. One study in a consulting firm found that humans were not able to discern AI writing well (100). Another study looked at so-called compelling disinformation, noting that humans were unable to distinguish between true and false tweets generated by GPT-3 and those written by real Twitter users (121). An additional study found that peer reviewers were not able to distinguish AI-generated from human-generated text in the journal article peer-review process for an applied linguistics journal (122). Finally, randomized-controlled experiments investigating novice and experienced teachers' ability to identify AI-generated texts showed that generative AI can write student essays in ways that were undetectable by teachers. Teachers were found to be overconfident in their source identification, and AI-generated essays were assessed more positively than student-written texts (123).

4.5. Real-World and Safe Use

Another broad concern is that most of the studies in the previous section showing impressive application of LLMs were conducted in simulated settings that may not reflect real-world use of generative AI. Indeed, there are few clinical trials that have assessed patient or healthcare delivery outcomes using all types of AI (124). There are also few assessments of the safety of generative AI, with some noting that legal liability for clinicians still rests with practitioners themselves when they use AI tools (125).

5. IMPACT OF GENERATIVE ARTIFICIAL INTELLIGENCE ON EDUCATION

We have already seen that the impact of generative AI on activities associated with education is profound and varied. This section describes recommendations for best practices going forward. Many thought leaders have expressed bold opinions on the use of AI in medical (and really all health professions) education. The dean of Harvard Medical School noted that "AI will and should change medical school" (126, p. 1820). Clinicians must be prepared to practice in a world of AI (127). Medical schools face the dual challenges of needing to teach about AI in practice but also adapt to its use by learners and faculty (128). Physicians must be prepared for the clinical algorithm era (129). Others conjecture that LLMs will change education and not destroy it, with one researcher noting that assessment may already be broken and stating, "if ChatGPT makes it easy to cheat on an assignment, teachers should throw out the assignment rather than ban the chatbot" (130). Another writer has asked, "Are we just grading robots?" (131). This section covers the following topics: competencies for the use of AI, prompt engineering, considerations for learning, role in clinical education, and roles outside clinical education.

5.1. Competencies for Use of Artificial Intelligence

Some work has focused on student competencies for AI, mainly use by students who will become practicing clinicians. One well-known set of competencies in clinical informatics was developed a decade ago (132) and was recently expanded to include AI (adapted from Reference 133):

- Find, search, and apply knowledge-based information to patient care and other clinical tasks
- Effectively read from, and write to, the electronic health record for patient care and other clinical activities
- Use and guide implementation of clinical decision support
- Provide care using population health management approaches
- Protect patient privacy and security
- Use information technology to improve patient safety
- Engage in quality measurement selection and improvement
- Use health information exchange to identify and access patient information across clinical settings
- Engage patients to improve their health and care delivery through personal health records and patient portals
- Maintain professionalism through use of information technology tools
- Provide clinical care via telemedicine and refer patients as indicated
- Apply personalized/precision medicine
- Participate in practice-based clinical and translational research
- Use and critique AI applications in clinical care

Others competency frameworks have been proposed that are more specific to AI. One focuses on the use of AI in primary care, proposing competencies in six domains (adapted from Reference 134):

- Foundational knowledge—what is this tool?
- Critical appraisal—should I use this tool?
- Medical decision making—when should I use this tool?
- Technical use—how do I use this tool?
- Patient communication—how should I communicate with patients regarding the use of the tool?
- Unintended consequences (cross-cutting)—what are the side effects of this tool?

Another framework, also with six areas, focuses on the use of AI-based tools by healthcare professionals more broadly (adapted from Reference 135):

- Basic knowledge of AI—explain what AI is and describe its healthcare applications
- Social and ethical implications of AI—explain how social, economic, and political systems influence AI-based tools and how these relationships impact justice, equity, and ethics
- AI-enhanced clinical encounters—carry out AI-enhanced clinical encounters that integrate diverse sources of information in creating patient-centered care plans
- Evidence-based evaluation of AI-based tools—evaluate the quality, accuracy, safety, contextual appropriateness, and biases of AI-based tools and their underlying datasets in providing care to patients and populations
- Workflow analysis for AI-based tools—analyze and adapt to changes in teams, roles, responsibilities, and workflows resulting from implementation of AI-based tools

- Practice-based learning and improvement regarding AI-based tools—participate in continuing professional development and practice-based improvement activities related to use of AI tools in healthcare

An additional framework presents a matrix of competencies for two types of clinicians, general and AI specialist, in three broad domains of health informatics, AI, and generative AI/LLMs (136).

5.2. Prompt Engineering

Another important competency in the era of generative AI is prompting, which is sometimes called prompt engineering. One primer focused on prompt engineering in the biomedical and health context and recommends (137):

- Be as specific as possible
- Describe the setting and provide the context around the question
- Experiment with different prompt styles
- Identify the overall goal of the prompt first
- Ask the LLM to play roles
- Iterate and refine prompts
- Use threads of prompts
- Ask open-ended questions
- Request examples
- Provide temporal awareness
- Set realistic expectations
- Use the one- or few-shot prompts
- Prompting LLMs for prompts

Another author notes a number of considerations for healthcare prompt engineering, raising issues that should be considered in the prompting process (138). These include recommendations similar to those above, but also some ethical and legal issues, such as not entering any personally identifiable information that may become part of the LLM; compliance with privacy standards, such as the Health Insurance Portability and Accountability Act and the General Data Protection Regulation; and aligning use with medical ethics.

5.3. Considerations for Learning

One concern for generative AI is the ease with which it can answer questions, pass tests, and substitute for deeper learning and understanding of a topic, all of which potentially undermine the ability of students to engage in critical thinking, both in their education and into their real-world professional work that follows their education. Although educators often try to make learning an enjoyable process, a recent meta-analysis found an association between mental effort and aversive affect, i.e., learning takes effort that is not always pleasant (139). This can lead to learners taking shortcuts, a process that started with use of Web search engines and now includes the use of generative AI. One recent study assessed 2,433 students over an 11-year period in 12 different college lecture courses (140). The researchers found that the percentage of students who did not show performance benefit by correctly answering homework questions increased from 14% in 2008 to 55% in 2017, presumably due to taking shortcuts to answering homework questions and not learning the underlying material as well. In the last 2 years of the study, students were asked how they completed their homework. Those who benefited from homework reported generating

their own answers, whereas students who reported copying the answers from other sources did not benefit from homework.

Other studies have specifically looked at the impact of ChatGPT in learning. One study assessed learning Python programming in a data science course and found that using LLMs as personal tutors by asking for explanations improved learning outcomes but that excessively asking LLMs to generate solutions impaired learning (141). The latter was made more adverse by allowing copy-and-paste availability. Another study evaluated learning math at the high school level (142). Both a plain version of ChatGPT-4 and one augmented with prompts aiming to provide more of a tutor mode were found to increase performance on exams by 48% and 127%, respectively. However, when the LLMs were taken away from students, subsequent scores averaged 17% worse than baseline for those using the plain version and a return to baseline for the tutor version. These studies show that LLMs may have adverse consequences on student learning, allowing easy answering of questions but not resulting in mastery of the material.

5.4. Role in Clinical Education

A number of authors have written about the role of generative AI in clinical education specifically. One narrative review looked at the potential of LLMs for medical education, noting potential advantages for students and faculty but also challenges (143). This narrative review noted advantages to students, such as more direct access to information, personalizing learning activities, and facilitating development of clinical skills. Potential benefits for faculty and instructors include the development of innovative approaches to pedagogy for complex medical concepts and facilitating student engagement. However, the review also noted a number of challenges, such as the potential for academic misconduct, students becoming overreliant on AI, dilution of critical thinking skills, concerns for the veracity and reliability of LLM-generated content, and the implications for all of these on teaching staff.

Another paper proposed a set of recommendations for medical faculty and their institutions (144). The authors recommended that educators increase their knowledge of AI and understand current approaches to its use in medical practice and education. This includes being familiar with strategies for successful AI integration into education and acting as stewards for its ethical use. Likewise, institutions must review and revise school policies, create new policies regarding the use of generative AI when necessary, support faculty development regarding AI, provide resources for teaching, and offer information-checking tools for originality and plagiarism to faculty.

An additional paper noted a number of use cases for how LLMs can be integrated into medical education (adapted from Reference 145):

- Practice generating differential diagnoses
- Streamlining the wide array of study resources to assist with devising a study plan
- Serving as a simulated patient or medical professor for interactive clinical cases
- Helping students review MCQs or generating new questions for additional practice
- Digesting lecture outlines and generating materials for flash cards
- Organizing information into tables to help build scaffolding for students to connect new information to previous knowledge

Other uses for LLMs in health professions education that have been put forth include creating checklists for common presentations and generating templates for common clinical scenarios (146), presenting potential patient problems to nurses and guiding students through clinical processes (147), and generating radiology board-style MCQs (148).

5.5. Roles Outside Clinical Education

Additional recommendations for the use of generative in AI in teaching have come from outside biomedicine and health. A couple of business school professors have written about “assigning AI,” noting the different ways it might be used (149):

- Mentor—providing feedback
- Tutor—direct instruction
- Coach—prompt metacognition
- Teammate—increase team performance
- Student—receive explanations
- Simulator—deliberate practice
- Tool—accomplish tasks

One of these authors has written a book about AI-human “co-intelligence” (150) and posted a library of prompts for use by educators (151). The two have also promoted the notion of instructors as innovators enabling novel forms of practice and application including simulations, mentoring, coaching, and co-creation (152). A number of other authors have written books describing approaches to using AI in education. One of these books covers topics such as policies and cheating, feedback and roleplaying, assignments, writing, and assessments (153). An additional business school professor described his approach to using LLMs to teach management sciences, noting that they should be used in writing assignments in an iterative manner for the duration of a course and that multiple-choice, brief response, and fill-in-the-blank tests should be avoided (154).

Clearly education at all levels, especially higher education, will need to adapt to AI. Another author notes that colleges and universities will need to move from “foundational to multifaceted” AI education, leverage experiential learning to use AI as well as surmount it, and use lifelong learning for reinvention, not just to acquire new skills (155).

6. CONCLUSIONS

Generative AI is having profound impact on biomedical and health professions and their education. It has been rapidly taken up by clinicians, students, and indeed all of society. Generative AI has been found to perform as well as human experts on many but not all intellectual tasks. In biomedicine and health, this includes medical board exams, answering clinical questions, solving clinical cases and applying clinical reasoning, summarizing information, and performing well in academic courses. For many activities, it helps less experienced individuals more than experts, although in other situations it can lead all people astray. There are also a number of limitations of generative AI, such as answers including hallucinations and confabulations, the inability to cite references and provide attribution, answers based on biased data, and algorithms and humans being poor at discerning text from generative AI. In addition, many of the studies have been done in simulated settings and not real-world practice. Generative AI is likewise being used widely in education, and a number of authors have provided recommendations for its optimal use in pedagogy. Clearly there are challenges for its use, especially if it undermines the acquisition of knowledge and skills for professional activity. However, all students and faculty, in biomedicine and health and beyond, must have a thorough understanding of it and be competent in its use.

This author has witnessed a number of transformations in education, including the transition from slide rules to calculators while in high school in the 1970s and the emergence of Google and other Internet search engines as a young faculty in the 1990s. The transition to generative AI, however, is probably more profound, since the earlier transitions simplified the steps of

problem-solving and critical thinking, whereas generative AI potentially replaces them. Educators must develop new approaches to teaching and student assessment in the era of generative AI, while healthcare, informatics, and educational professionals must be competent with AI as much as any other tool in professional practice.

FUTURE ISSUES

1. Will the performance of large language models (LLMs) continue to improve in biomedical and health tasks, or will it level off?
2. How will the use of LLMs be optimized and validated beyond simulated use and implemented in the larger workflow of biomedical and health professional practice that includes clinical care, research, administration, and other areas?
3. What evidence will be required of assertions students and professionals obtain from generative artificial intelligence (AI)?
4. What will be the optimal policies for the use of LLMs in education?
5. How will we assess student learning when generative AI tools are readily available?
6. Knowing that in most biomedical and health disciplines, students can no longer memorize all the knowledge and skills of their disciplines, what core must be mastered to provide a foundation and perspective that will enable advanced thinking?
7. How will students and instructors minimize overreliance on LLMs that may undermine their broader professional competence?

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

1. Copeland BJ. 2024. *Encyclopedia Britannica*. Artificial intelligence. <https://www.britannica.com/technology/artificial-intelligence>
2. Khare Y. 2023. Generative AI versus predictive AI: What is the difference? *Analytics Vidhya*. <https://www.analyticsvidhya.com/blog/2023/09/generative-ai-vs-predictive-ai>
3. McCarthy J, Feigenbaum EA. 1990. In memoriam: Arthur Samuel: pioneer in machine learning. *AI Mag*. 11(3):10
4. Shah C. 2022. *A Hands-On Introduction to Machine Learning*. Cambridge, UK: Cambridge Univ. Press
5. Lambert N, Castricato L, von Werra L, Havrilla A. 2022. Illustrating reinforcement learning from human feedback (RLHF). *Hugging Face*. <https://huggingface.co/blog/rlhf>
6. Bahree A. 2024. *Generative AI in Action*. Shelter Island, NY: Manning
7. Shao M, Basit A, Karri R, Shafique M. 2024. Survey of different large language model architectures: trends, benchmarks, and challenges. *IEEE Access* 12:188664–706
8. Omiye JA, Gui H, Rezaei SJ, Daneshjou R. 2024. Large language models in medicine: the potentials and pitfalls: a narrative review. *Ann. Intern. Med.* 177(2):210–20. Erratum. 2024. *Ann. Intern. Med.* 177:548
9. Shanmugam D, Agrawal M, Movva R, Chen IY, Ghassemi M, et al. 2024. Generative AI in medicine. *arXiv:2412.10337 [cs.LG]*
10. Xiao H, Zhou F, Liu X, Liu T, Li Z, et al. 2024. A comprehensive survey of large language models and multimodal large language models in medicine. *Inform. Fusion* 117:102888



11. Nam J. 2023. 56% of college students have used AI on assignments or exams. *BestColleges*. <https://www.bestcolleges.com/research/most-college-students-have-used-ai-survey>
12. Impact Res. *AI chatbots in schools: findings from a poll of K-12 teachers, students, parents, and college undergraduates*. Rep., Impact Res., Washington, DC. <https://8ce82b94a8c4fdc3ea6d-b1d233e3bc3cb10858bea65ff05e18f2.ssl.cf2.rackcdn.com/bf/24/cd3646584af89e7c668c7705a006/deck-impact-analysis-national-schools-tech-tracker-may-2024-1.pdf>
13. Bick A, Blandin A, Deming DJ. 2024. *The rapid adoption of generative AI*. NBER Work. Pap. 32966
14. Blease CR, Locher C, Gaab J, Hägglund M, Mandl KD. 2024. Generative artificial intelligence in primary care: an online survey of UK general practitioners. *BMJ Health Care Inform*. 31(1):e101102
15. Presiado M, Montero A, Lopes L. 2024. KFF health misinformation tracking poll: artificial intelligence and health information. *KFF*. <https://www.kff.org/health-misinformation-and-trust/poll-finding/kff-health-misinformation-tracking-poll-artificial-intelligence-and-health-information>
16. Jin D, Pan E, Oufattole N, Weng WH, Fang H, Szolovits P. 2021. What disease does this patient have? A large-scale open domain question answering dataset from medical exams. *Appl. Sci*. 11(14):6421
17. Kung TH, Cheatham M, Medenilla A, Sillos C, De Leon L, et al. 2023. Performance of ChatGPT on USMLE: potential for AI-assisted medical education using large language models. *PLOS Digit. Health*. 2(2):e0000198
18. Nori H, King N, McKinney SM, Carignan D, Horvitz E. 2023. Capabilities of GPT-4 on medical challenge problems. arXiv:2303.13375 [cs.CL]
19. Corrado G, Barral J. 2024. Advancing medical AI with Med-Gemini. *Google Research*. <http://research.google/blog/advancing-medical-ai-with-med-gemini>
20. Horvitz E, Nori H, Usuyama N. 2024. Advances in run-time strategies in foundation models. *Microsoft Research Blog*. <https://www.microsoft.com/en-us/research/blog/advances-in-run-time-strategies-for-next-generation-foundation-models>
21. Nori H, Usuyama N, King N, McKinney SM, Fernandes X, et al. 2024. From Medprompt to o1: exploration of run-time strategies for medical challenge problems and beyond. arXiv:2411.03590 [cs.CL]
22. Brin D, Sorin V, Vaid A, Soroush A, Glicksberg BS, et al. 2023. Comparing ChatGPT and GPT-4 performance in USMLE soft skill assessments. *Sci. Rep*. 13(1):16492
23. Adams LC, Truhn D, Busch F, Dorfner F, Nawabi J, et al. 2024. Llama 3 challenges proprietary state-of-the-art large language models in radiology board-style examination questions. *Radiology* 312(2):e241191
24. Bhayana R, Bleakney RR, Krishna S. 2023. GPT-4 in radiology: improvements in advanced reasoning. *Radiology* 307(5):e230987
25. Ali R, Tang OY, Connolly ID, Zadnik Sullivan PL, Shin JH, et al. 2023. Performance of ChatGPT and GPT-4 on neurosurgery written board examinations. *Neurosurgery* 93(6):1353–65
26. Kumah-Crystal Y, Mankowitz S, Embi P, Lehmann CU. 2023. ChatGPT and the Clinical Informatics Board Examination: The end of unproctored maintenance of certification? *J. Am. Med. Inform. Assoc*. 30(9):1558–60
27. Katz U, Cohen E, Shachar E, Somer J, Fink A, et al. 2024. GPT versus resident physicians—a benchmark based on official board scores. *NEJM AI* 1(5):A1dbp2300192
28. Li SW, Kemp MW, Logan SJS, Dimri PS, Singh N, et al. 2023. ChatGPT outscored human candidates in a virtual objective structured clinical examination in obstetrics and gynecology. *Am. J. Obstet. Gynecol*. 229(2):172.e1–12
29. Teebagy S, Colwell L, Wood E, Yaghy A, Faustina M. 2023. Improved performance of ChatGPT-4 on the OKAP exam: a comparative study with ChatGPT-3.5. medRxiv 2023.04.03.23287957. <https://doi.org/10.1101/2023.04.03.23287957>
30. Wu S, Koo M, Blum L, Black A, Kao L, et al. 2024. Benchmarking open-source large language models, GPT-4 and Claude 2 on multiple-choice questions in nephrology. *NEJM AI* 1(2):A1dbp2300092
31. Hanna RE, Smith LR, Mhaskar R, Hanna K. 2024. Performance of language models on the family medicine in-training exam. *Fam. Med*. 56(9):555–60
32. Wang T, Mainous AG, Stelter K, O'Neill TR, Newton WP. 2024. Performance evaluation of the Generative Pre-trained Transformer (GPT-4) on the Family Medicine In-Training Examination. *J. Am. Board Fam. Med*. 37(4):528–82

33. Payne DL, Purohit K, Borrero WM, Chung K, Hao M, et al. 2024. Performance of GPT-4 on the American College of Radiology In-Training Examination: evaluating accuracy, model drift, and fine-tuning. *Acad. Radiol.* 31(7):3046–54
34. Chen S, Kann BH, Foote MB, Aerts HJWL, Savova GK, et al. 2023. Use of artificial intelligence chatbots for cancer treatment information. *JAMA Oncol.* 9(10):1459–62
35. Rydzewski NR, Dinakaran D, Zhao SG, Ruppel E, Turkbey B, et al. 2024. Comparative evaluation of LLMs in clinical oncology. *NEJM AI* 1(5):AIoa2300151
36. Ferber D, Wiest IC, Wölflin G, Ebert MP, Beutel G, et al. 2024. GPT-4 for information retrieval and comparison of medical oncology guidelines. *NEJM AI* 1(6):AIcs2300235
37. Goodman RS, Patrinely JR, Stone CA, Zimmerman E, Donald RR, et al. 2023. Accuracy and reliability of chatbot responses to physician questions. *JAMA Netw. Open.* 6(10):e2336483
38. Duong D, Solomon BD. 2023. Analysis of large-language model versus human performance for genetics questions. *Eur. J. Hum. Genet.* 32(4):466–68
39. Dash D, Thapa R, Banda JM, Swaminathan A, Cheatham M, et al. 2023. Evaluation of GPT-3.5 and GPT-4 for supporting real-world information needs in healthcare delivery. arXiv:2304.13714 [cs.AI]
40. Rodman A, Buckley TA, Manrai AK, Morgan DJ. 2023. Artificial intelligence versus clinician performance in estimating probabilities of diagnoses before and after testing. *JAMA Netw. Open.* 6(12):e2347075
41. Zakka C, Shad R, Chaurasia A, Dalal AR, Kim JL, et al. 2024. Almanac—retrieval-augmented language models for clinical medicine. *NEJM AI* 1(2):AIoa2300068
42. Low YS, Jackson ML, Hyde RJ, Brown RE, Sanghavi NM, et al. 2024. Answering real-world clinical questions using large language model based systems. arXiv:2407.00541 [cs.CL]
43. Xie Y, Wu J, Tu H, Yang S, Zhao B, et al. 2024. A preliminary study of o1 in medicine: Are we closer to an AI doctor? arXiv:2409.15277 [cs.CL]
44. Semigran HL, Levine DM, Nundy S, Mehrotra A. 2016. Comparison of physician and computer diagnostic accuracy. *JAMA Intern. Med.* 176(12):1860–1
45. Benoit JRA. 2023. ChatGPT for clinical vignette generation, revision, and evaluation. medRxiv 2023.02.04.23285478. <https://doi.org/10.1101/2023.02.04.23285478>
46. Levine DM, Tuwani R, Kompa B, Varma A, Finlayson SG, et al. 2024. The diagnostic and triage accuracy of the GPT-3 artificial intelligence model: an observational study. *Lancet Digit. Health* 6(8):e555–61
47. Kanjee Z, Crowe B, Rodman A. 2023. Accuracy of a generative artificial intelligence model in a complex diagnostic challenge. *JAMA* 330(1):78–80
48. Eriksen AV, Möller S, Ryg J. 2023. Use of GPT-4 to diagnose complex clinical cases. *NEJM AI* 1(1):AIp2300031
49. McDuff D, Schaekermann M, Tu T, Palepu A, Wang A, et al. 2023. Towards accurate differential diagnosis with large language models. arXiv:2312.00164 [cs.CY]
50. Rao A, Pang M, Kim J, Kamineni M, Lie W, et al. 2023. Assessing the utility of ChatGPT throughout the entire clinical workflow: development and usability study. *J. Med. Internet Res.* 25:e48659
51. ten Berg H, van Bakel B, van de Wouw L, Jie KE, Schipper A, et al. 2024. ChatGPT and generating a differential diagnosis early in an emergency department presentation. *Ann. Emerg. Med.* 83(1):83–86
52. Levkovich I, Elyoseph Z. 2023. Identifying depression and its determinants upon initiating treatment: ChatGPT versus primary care physicians. *Fam. Med. Commun. Health* 11(4):e002391
53. Iqbal U, Lee LTJ, Rahmanti AR, Celi LA, Li YCJ. 2024. Can large language models provide secondary reliable opinion on treatment options for dermatological diseases? *J. Am. Med. Inform. Assoc.* 31(6):1341–47
54. Chen A, Chen DO, Tian L. 2023. Benchmarking the symptom-checking capabilities of ChatGPT for a broad range of diseases. *J. Am. Med. Inform. Assoc.* 31(9):2084–88
55. Cabral S, Restrepo D, Kanjee Z, Wilson P, Crowe B, et al. 2024. Clinical reasoning of a generative artificial intelligence model compared with physicians. *JAMA Intern. Med.* 184(5):581–83
56. Berner ES, Webster GD, Shugerman AA, Jackson JR, Algina J, et al. 1994. Performance of four computer-based diagnostic systems. *N. Engl. J. Med.* 330(25):1792–96



57. Goh E, Gallo R, Hom J, Strong E, Weng Y, et al. 2024. Influence of a large language model on diagnostic reasoning: a randomized clinical vignette study. medRxiv 2024.03.12.24303785. <https://doi.org/10.1101/2024.03.12.24303785>
58. Goh E, Gallo R, Strong E, Weng Y, Kerman H, et al. 2024. Large language model influence on management reasoning: a randomized controlled trial. medRxiv 2024.08.05.24311485. <https://doi.org/10.1101/2024.08.05.24311485>
59. Hager P, Jungmann F, Holland R, Bhagat K, Hubrecht I, et al. 2024. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nat. Med.* 30(9):2613–22
60. Tu T, Palepu A, Schaekermann M, Saab K, Freyberg J, et al. 2024. Towards conversational diagnostic AI. arXiv:2401.05654 [cs.AI]
61. O'Sullivan JW, Palepu A, Saab K, Weng WH, Cheng Yet al. 2024. Towards democratization of subspecialty medical expertise. arXiv:2410.03741 [cs.HC]
62. Arvidsson R, Gunnarsson R, Entezarjou A, Sundemo D, Wikberg C. 2024. ChatGPT (GPT-4) versus doctors on complex cases of the Swedish family medicine specialist examination: an observational comparative study. *BMJ Open* 14(12):e086148
63. Johri S, Jeong J, Tran BA, Schlessinger DI, Wongvibulsin S, et al. 2025. An evaluation framework for clinical use of large language models in patient interaction tasks. *Nat. Med.* 31:77–86
64. Han T, Adams LC, Bressem KK, Busch F, Nebelung S, Truhn D. 2024. Comparative analysis of multimodal large language model performance on clinical vignette questions. *JAMA* 331(15):1320–1
65. Jin Q, Chen F, Zhou Y, Xu Z, Cheung JM, et al. 2024. Hidden flaws behind expert-level accuracy of multimodal GPT-4 vision in medicine. *NPJ Digit. Med.* 7(1):190
66. Xu S, Yang L, Kelly C, Sieniek M, Kohlberger T, et al. 2023. ELIXR: towards a general purpose X-ray artificial intelligence system through alignment of large language models and radiology vision encoders. arXiv:2308.01317 [cs.CV]
67. Zhou HY, Yu Y, Wang C, Zhang S, Gao Y, et al. 2023. A transformer-based representation-learning model with unified processing of multimodal input for clinical diagnostics. *Nat. Biomed. Eng.* 7(6):743–55
68. Huang J, Neill L, Wittbrodt M, Melnick D, Klug M, et al. 2023. Generative artificial intelligence for chest radiograph interpretation in the emergency department. *JAMA Netw. Open* 6(10):e2336100
69. Li C, Wong C, Zhang S, Usuyama N, Liu H, et al. 2023. LLaVA-Med: training a large language-and-vision assistant for biomedicine in one day. arXiv:2306.00890 [cs.CV]
70. Bala W, Li H, Moon J, Trivedi H, Gichoya J, Balthazar P. 2024. Enhancing radiology training with GPT-4: pilot analysis of automated feedback in trainee preliminary reports. *Curr. Probl. Diagn. Radiol.* In press. <https://doi.org/10.1067/j.cpradiol.2024.08.003>
71. Gertz RJ, Dratsch T, Bunck AC, Lennartz S, Iuga AI, et al. 2024. Potential of GPT-4 for detecting errors in radiology reports: implications for reporting accuracy. *Radiology* 311(1):e232714
72. Liang W, Zhang Y, Cao H, Wang B, Ding DY, et al. 2024. Can large language models provide useful feedback on research papers? A large-scale empirical analysis. *NEJM AI* 1(8):AIoa2400196
73. Hake J, Crowley M, Coy A, Shanks D, Eoff A, et al. 2024. Quality, accuracy, and bias in ChatGPT-based summarization of medical abstracts. *Ann. Fam. Med.* 22(2):113–20
74. Skarlinski MD, Cox S, Laurent JM, Braza JD, Hinks M, et al. 2024. Language agents achieve superhuman synthesis of scientific knowledge. arXiv:2409.13740 [cs.CL]
75. Shyr C, Grout RW, Kennedy N, Akdas Y, Tischbein M, et al. 2024. Leveraging artificial intelligence to summarize abstracts in lay language for increasing research accessibility and transparency. *J. Am. Med. Inform. Assoc.* 31(10):2294–303
76. Han C, Kim DW, Kim S, You SC, Park JY, et al. 2023. *Evaluation of GPT-4 for 10-year cardiovascular risk prediction: insights from the UK Biobank and KoGES data*. SSRN Work. Pap. 4583995. <https://papers.ssrn.com/abstract=4583995>
77. Williams CYK, Zack T, Miao BY, Sushil M, Wang M, et al. 2024. Use of a large language model to assess clinical acuity of adults in the emergency department. *JAMA Netw. Open* 7(5):e248895
78. Glicksberg BS, Timsina P, Patel D, Sawant A, Vaid A, et al. 2024. Evaluating the accuracy of a state-of-the-art large language model for prediction of admissions from the emergency room. *J. Am. Med. Inform. Assoc.* 31:1921–28

79. Zelin C, Chung WK, Jeanne M, Zhang G, Weng C. 2024. Rare disease diagnosis using knowledge guided retrieval augmentation for ChatGPT. *J. Biomed. Inform.* 157:104702
80. Stribling D, Xia Y, Amer MK, Graim KS, Mulligan CJ, Renne R. 2024. The model student: GPT-4 performance on graduate biomedical science exams. *Sci. Rep.* 14(1):5670
81. Hersh W, Fultz Hollis K. 2024. Results and implications for generative AI in a large introductory biomedical and health informatics course. *NPJ Digit. Med.* 7(1):247
82. Avramovic S, Avramovic I, Wojtusiak J. 2024. Exploring the impact of GitHub Copilot on health informatics education. *Appl. Clin. Inform.* 15(5):1121–29
83. Grattafiori A, Dubey A, Jauhri A, Pandey A, Kadian A, et al. 2024. The Llama 3 herd of models. arXiv:2407.21783 [cs.AI]
84. Scarfe P, Watcham K, Clarke A, Roesch E. 2024. A real-world test of artificial intelligence infiltration of a university examinations system: a “Turing test” case study. *PLOS ONE* 19(6):e0305354
85. Rein D, Hou BL, Stickland AC, Petty J, Pang RY, et al. GPQA: a graduate-level Google-proof Q&A benchmark. arXiv:2311.12022 [cs.AI]
86. Jones N. 2024. ‘In awe’: scientists impressed by latest ChatGPT model o1. *Nature* 634:275–76
87. Denny P, Prather J, Becker BA, Finnie-Ansley J, Hellas A, et al. 2024. Computing education in the era of generative AI. *Commun. ACM* 67(2):56–67
88. Poldrack RA, Lu T, Beguš G. 2023. AI-assisted coding: experiments with GPT-4. arXiv:2304.13187 [cs.AI]
89. Johnson M. 2024. Generative AI and CS education. *Commun. ACM* 67(4):23–24
90. Choi JH, Monahan A, Schwarcz D. *Lawyering in the age of artificial intelligence*. SSRN Work. Pap. 4626276. <https://papers.ssrn.com/abstract=4626276>
91. Cheng L, Li X, Bing L. 2023. Is GPT-4 a good data analyst? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, ed. H Bouamor, J Pino, K Bali, pp. 9496–514. Kerrville, TX: Assoc. Comput. Linguist.
92. Hong S, Lin Y, Liu B, Liu B, Wu B, et al. 2024. Data Interpreter: an LLM agent for data science. arXiv:2402.18679 [cs.AI]
93. Bubeck S, Chandrasekaran V, Eldan R, Gehrke J, Horvitz E, et al. arXiv.org. 2023. Sparks of artificial general intelligence: early experiments with GPT-4. arXiv:2303.12712 [cs.CL]
94. Koivisto M, Grassini S. 2023. Best humans still outperform artificial intelligence in a creative divergent thinking task. *Sci. Rep.* 13(1):13601
95. Bang Y, Cahyawijaya S, Lee N, Dai W, Su D, et al. 2023. A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity. arXiv:2302.04023 [cs.CL]
96. Noy S, Zhang W. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654):187–92
97. Kim A, Muhn M, Nikolaev VV. 2024. *Financial statement analysis with large language models*. SSRN Work. Pap. 4835311. <https://papers.ssrn.com/abstract=4835311>
98. Shaffer M, Wang CCY. 2024. *Scaling core earnings measurement with large language models*. SSRN Work. Pap. 4979501. <https://papers.ssrn.com/abstract=4979501>
99. Cui Z (K), Demirer M, Jaffe S, Musolff L, Peng S, Salz T. 2024. *The effects of generative AI on high skilled work: evidence from three field experiments with software developers*. SSRN Work. Pap. 4945566. <https://papers.ssrn.com/abstract=4945566>
100. Dell’Acqua F, McFowland E, Mollick ER, Lifshitz-Assaf H, Kellogg K, et al. 2023. *Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality*. SSRN Work. Pap. 4573321. <https://papers.ssrn.com/abstract=4573321>
101. Norlen N, Barrett G. 2023. The Dictionary.com word of the year is *hallucinate*. *Dictionary.com*. <https://content.dictionary.com/word-of-the-year-2023>
102. Augenstein I, Baldwin T, Cha M, Chakraborty T, Ciampaglia GL, et al. 2023. Factuality challenges in the era of large language models. arXiv:2310.05189 [cs.CL]
103. Walters WH, Wilder EI. 2023. Fabrication and errors in the bibliographic citations generated by ChatGPT. *Sci. Rep.* 13(1):14045
104. Wu K, Wu E, Cassasola A, Zhang A, Wei K, et al. 2024. How well do LLMs cite relevant medical references? An evaluation framework and analyses. arXiv:2402.02008 [cs.CL]



105. Hersh W. 2024. Search still matters: information retrieval in the era of generative AI. *J. Am. Med. Inform. Assoc.* 31(9):2159–61
106. Omiye JA, Lester JC, Spichak S, Rotemberg V, Daneshjou R. 2023. Large language models propagate race-based medicine. *NPJ Digit. Med.* 6(1):195
107. Zack T, Lehman E, Suzgun M, Rodriguez JA, Celi LA, et al. 2024. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *Lancet Digit. Health* 6(1):e12–22
108. Fang X, Che S, Mao M, Zhang H, Zhao M, Zhao X. 2024. Bias of AI-generated content: an examination of news produced by large language models. *Sci. Rep.* 14(1):5224
109. Kapoor S, Narayanan A. 2023. Leakage and the reproducibility crisis in machine-learning-based science. *Patterns* 4(9):100804
110. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, et al. 2021. The clinician and dataset shift in artificial intelligence. *N. Engl. J. Med.* 385(3):283–86
111. Knaus WA. 2022. Prognostic modeling and major dataset shifts during the COVID-19 pandemic: What have we learned for the next pandemic? *JAMA Health Forum* 3(5):e221103
112. Aronson SJ, Machini K, Shin J, Sriraman P, Hamill S, et al. 2024. GPT-4 performance, nondeterminism, and drift in genetic literature review. *NEJM AI* 1(9):AIcs2400245
113. Tang R, Chuang YN, Hu X. 2024. The science of detecting LLM-generated text. *Commun. ACM* 67(4):50–59
114. Gao CA, Howard FM, Markov NS, Dyer EC, Ramesh S, et al. 2023. Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers. *NPJ Digit. Med.* 6(1):75
115. Desaire H, Chua AE, Isom M, Jarosova R, Hua D. 2023. Distinguishing academic science writing from humans or ChatGPT with over 99% accuracy using off-the-shelf machine learning tools. *Cell Rep. Phys. Sci.* 4(6):101426
116. Desaire H, Chua AE, Kim MG, Hua D. 2023. Accurately detecting AI text when ChatGPT is told to write like a chemist. *Cell Rep. Phys. Sci.* 4(11):101672
117. Sadasivan VS, Kumar A, Balasubramanian S, Wang W, Feizi S. 2023. Can AI-generated text be reliably detected? arXiv:2303.11156 [cs.CL]
118. Odri G-A, Yun Yoon DJ. 2023. Detecting generative artificial intelligence in scientific articles: evasion techniques and implications for scientific integrity. *Orthop. Traumatol. Surg. Res.* 109(8):103706
119. Perkins M, Roe J, Vu BH, Postma D, Hickerson D, et al. 2024. GenAI detection tools, adversarial techniques and implications for inclusivity in higher education. arXiv:2403.19148 [cs.CY]
120. Liang W, Yuksekgonul M, Mao Y, Wu E, Zou J. 2023. GPT detectors are biased against non-native English writers. *Patterns* 4(7):100779
121. Spitale G, Biller-Andorno N, Germani F. 2023. AI model GPT-3 (dis)informs us better than humans. *Sci. Adv.* 9(26):eadh1850
122. Casal JE, Kessler M. 2023. Can linguists distinguish between ChatGPT/AI and human writing?: A study of research ethics and academic publishing. *Res. Methods Appl. Linguistics* 2(3):100068
123. Fleckenstein J, Meyer J, Jansen T, Keller SD, Köller O, Möller J. 2024. Do teachers spot AI? Evaluating the detectability of AI-generated texts among student essays. *Comput. Educ. Artif. Intell.* 6:100209
124. Hersh W. 2024. Translational AI: a necessity and opportunity for biomedical informatics and data science. *NLM Musings from the Mezzanine*. <https://nlmdirector.nlm.nih.gov/2024/02/07/translational-ai-a-necessity-and-opportunity-for-biomedical-informatics-and-data-science>
125. Gliadkovskaya A. 2024. Some doctors are using ChatGPT to assist with clinical decisions. Is it safe? *Fierce Healthcare*. <https://www.fiercehealthcare.com/special-reports/some-doctors-are-using-public-generative-ai-tools-chatgpt-clinical-decisions-it>
126. Hsuen Y, Abbasi J. 2023. AI will—and should—change medical school, says Harvard's dean for medical education. *JAMA* 330(19):1820–23
127. James CA, Wachter RM, Woolliscroft JO. 2022. Preparing clinicians for a clinical world influenced by artificial intelligence. *JAMA* 327(14):1333–34
128. Cooper A, Rodman A. 2023. AI and medical education—a 21st-century Pandora's box. *N. Engl. J. Med.* 389:385–87



129. Goodman KE, Rodman AM, Morgan DJ. 2023. Preparing physicians for the clinical algorithm era. *N. Engl. J. Med.* 389(6):483–87
130. Heaven WD. 2023. ChatGPT is going to change education, not destroy it. *MIT Technology Review*. <https://www.technologyreview.com/2023/04/06/1071059/chatgpt-change-not-destroy-education-openai>
131. McMurtrie B. 2024. Professors ask: Are we just grading robots? *The Chronicle of Higher Education*. <https://www.chronicle.com/article/professors-ask-are-we-just-grading-robots>
132. Hersh WR, Gorman PN, Biagioli FE, Mohan V, Gold JA, Mejicano GC. 2014. Beyond information retrieval and electronic health record use: competencies in clinical informatics for medical education. *Adv. Med. Educ. Pract.* 5:205–12
133. Hersh W. 2023. Physician and medical student competence in AI must include broader competence in clinical informatics. *Informatics Professor*. <https://informaticsprofessor.blogspot.com/2023/09/physician-and-medical-student.html>
134. Liaw W, Kueper JK, Lin S, Bazemore A, Kakadiaris I. 2022. Competencies for the use of artificial intelligence in primary care. *Ann. Fam. Med.* 20(6):559–63
135. Russell RG, Lovett Novak L, Patel M, Garvey KV, Craig KJT, et al. 2023. Competencies for the use of artificial intelligence-based tools by health care professionals. *Acad. Med.* 98(3):348–56
136. McCoy LG, Ci Ng FY, Sauer CM, Yap Legaspi KE, Jain B, et al. 2024. Understanding and training for the impact of large language models and artificial intelligence in healthcare practice: a narrative review. *BMC Med. Educ.* 24(1):1096
137. Meskó B. 2023. Prompt engineering as an important emerging skill for medical professionals: tutorial. *J. Med. Internet Res.* 25:e50638
138. Rajaratnam V. 2023. Guide for healthcare prompt engineering with ChatGPT. *LinkedIn*. <https://www.linkedin.com/pulse/guide-healthcare-prompt-engineering-chatgpt-vaikunthan-rajaratnam>
139. David L, Vassena E, Bijleveld E. 2024. The unpleasantness of thinking: a meta-analytic review of the association between mental effort and negative affect. *Psychol. Bull.* 150(9):1070–93
140. Glass AL, Kang M. 2022. Fewer students are benefiting from doing their homework: an eleven-year study. *Educ. Psychol.* 42(2):185–99
141. Lehmann M, Cornelius PB, Sting FJ. 2024. *AI meets the classroom: When does ChatGPT harm learning?* SSRN Work. Pap. 4941259. <https://papers.ssrn.com/abstract=4941259>
142. Bastani H, Bastani O, Sungu A, Ge H, Kabakçı Ö, Mariman R. 2024. *Generative AI can harm learning*. SSRN Work. Pap. 4895486. <https://papers.ssrn.com/abstract=4895486>
143. Benítez TM, Xu Y, Boudreau JD, Kow AWC, Bello F, et al. 2024. Harnessing the potential of large language models in medical education: promise and pitfalls. *J. Am. Med. Inform. Assoc.* 31(3):776–83
144. Boscardin CK, Gin B, Golde PB, Hauer KE. 2023. ChatGPT and generative artificial intelligence for medical education: potential impact and opportunity. *Acad. Med.* 99(1):22–27
145. Safranek CW, Sidamon-Eristoff AE, Gilson A, Chartash D. 2023. The role of large language models in medical education: applications and implications. *JMIR Med. Educ.* 9:e50945
146. Bair H, Djulbegovic M, Taylor Gonzalez D. 2023. We use ChatGPT in medical practice – and you can too. *MedPageToday*. <https://www.medpagetoday.com/opinion/second-opinions/106085>
147. Gosak L, Pruinelli L, Topaz M, Štiglic G. 2024. The ChatGPT effect and transforming nursing education with generative AI: discussion paper. *Nurse Educ. Pract.* 75:103888
148. Mistry NP, Saeed H, Rafique S, Le T, Obaid H, Adams SJ. 2024. Large language models as tools to generate radiology board-style multiple-choice questions. *Acad. Radiol.* 31(9):3872–78
149. Mollick ER, Mollick L. *Assigning AI: seven approaches for students, with prompts*. SSRN Work. Pap. 4475995. <https://papers.ssrn.com/abstract=4475995>
150. Mollick E. 2024. *Co-Intelligence: Living and Working with AI*. New York: Portfolio
151. Mollick E. Prompt library. *More Useful Things: AI Resources*. <https://www.moreusefulthings.com/prompts>
152. Mollick E, Mollick L. 2024. How to use AI to create role-play scenarios for your students. *Harvard Business Publishing Education*. <https://hbsp.harvard.edu/inspiring-minds/using-generative-ai-to-create-role-play-scenarios-for-students>



153. Bowen JA. 2024. *Teaching with AI: A Practical Guide to a New Era of Human Learning*. Baltimore, MD: Johns Hopkins Univ. Press
154. McMurtrie B. 2024. Getting in Sync with Students About AI. *The Chronicle of Higher Education*. <https://www.chronicle.com/newsletter/teaching/2024-07-05>
155. Aoun JE. 2024. How higher ed can adapt to the challenges of AI. *The Chronicle of Higher Education*. <https://www.chronicle.com/article/how-higher-ed-can-adapt-to-the-challenges-of-ai>