

ISPOR Report

Generative Artificial Intelligence for Health Technology Assessment: Opportunities, Challenges, and Policy Considerations: An ISPOR Working Group Report

Rachael L. Fleurence, PhD, Jiang Bian, PhD, Xiaoyan Wang, PhD, Hua Xu, PhD, Dalia Dawoud, PhD, Mitchell Higashi, PhD, Jagpreet Chhatwal, PhD, on behalf of the ISPOR Working Group on Generative AI

A B S T R A C T

Objectives: To provide an introduction to the uses of generative artificial intelligence (AI) and foundation models, including large language models, in the field of health technology assessment (HTA).

Methods: We reviewed applications of generative AI in 3 areas: systematic literature reviews, real-world evidence, and health economic modeling.

Results: (1) Literature reviews: generative AI has the potential to assist in automating aspects of systematic literature reviews by proposing search terms, screening abstracts, extracting data, and generating code for meta-analyses; (2) real-world evidence: generative AI can facilitate automating processes and analyze large collections of real-world data, including unstructured clinical notes and imaging; (3) health economic modeling: generative AI can aid in the development of health economic models, from conceptualization to validation. Limitations in the use of foundation models and large language models include challenges surrounding their scientific rigor and reliability, the potential for bias, implications for equity, as well as nontrivial concerns regarding adherence to regulatory and ethical standards, particularly in terms of data privacy and security. Additionally, we survey the current policy landscape and provide suggestions for HTA agencies on responsibly integrating generative AI into their workflows, emphasizing the importance of human oversight and the fast-evolving nature of these tools.

Conclusions: Although generative AI technology holds promise with respect to HTA applications, it is still undergoing rapid developments and improvements. Continued careful evaluation of their applications to HTA is required. Both developers and users of research incorporating these tools, should familiarize themselves with their current capabilities and limitations.

Keywords: artificial intelligence, economic modeling methods, generative AI, large language models, real world evidence, systematic reviews.

VALUE HEALTH. 2025; 28(2):175–183

Highlights

- Generative artificial intelligence (AI) and foundation models, including large language models, have the potential to transform health technology assessment (HTA) by automating elements of evidence synthesis, enhancing real-world evidence generation, and streamlining health economic modeling.
- Despite their potential, current generative AI applications are in their early stages and present limitations, including issues of scientific validity and reliability, risk of bias and impact on equity, and regulatory and ethical considerations. These will continue to require careful evaluation and human oversight in their application to HTA for the foreseeable future.
- The article reviews the current policy landscape of generative AI, and the authors offer suggestions for HTA agencies, emphasizing the need for integrating the use of foundation models into guidance, harmonizing standards, and investing in training to responsibly integrate generative AI into their assessment processes.

Introduction

Artificial intelligence (AI) has been defined as “the science and engineering of making intelligent machines.”¹ Advancements in AI and especially machine learning (ML) have been rapidly progressing for several decades.^{2,3} However, the recent emergence of generative AI technologies, such as large language models (LLMs), has ushered in a transformative era across numerous fields, including science, medicine, and other areas of human activity^{4–6} (Box 1) (Fig. 1).

Generative AI uses sophisticated ML models, particularly a class of deep neural networks, which can generate language, images, and code in response to free-text prompts provided by users. These models are trained using self-supervised learning techniques on vast amounts of existing data, allowing them to identify patterns and relationships in the underlying data autonomously. Their

training requires substantial computational power, often relying on specialized hardware, especially graphics processing units.⁷

In the field of Health Economics and Outcomes Research (HEOR), generative AI has the potential to transform evidence generation methods used in health technology assessments (HTAs). Despite this promise, the exploration of opportunities, limitations, and risks associated with developing, evaluating, and deploying these models for supporting different approaches to evidence generation remains sparse, given the recency of these new tools. This article seeks to fill this gap by offering a review of the landscape of generative AI's applications in methodological areas relevant to HTA development.

The perspective taken in this article is that, in this current phase, foundation models should be used to augment human activities for which humans remain fully accountable. Researchers are ultimately responsible for the quality and accuracy of results generated with the help of foundation models and for using existing guidelines and checklists to report results. Nevertheless, foundation models do add an additional layer of complexity that needs to be acknowledged and addressed during the analysis and the reporting.

The immediate audience of the article is HEOR professionals, but we anticipate that it will also be useful to other stakeholders, particularly HTA evidence developers and users, regulators, and manufacturers. The article is structured as follows. We begin with a concise history of AI and overview of generative AI. Next, we examine the opportunities and challenges of utilizing generative AI to support significant aspects of HEOR. This includes (1) literature reviews and evidence synthesis, (2) real-world evidence (RWE), and (3) health economic modeling. We continue with a description of the current limitations and risks associated with generative AI tools. Next, we provide a summary overview of the current policy landscape on AI. We conclude with the results of a brief survey of HTA agencies gauging their current use of tools with generative AI and provide suggestions to HTA agencies on preparing for a more widespread adoption of generative AI tools.

A Brief History of AI and Generative AI

AI has evolved significantly since its conception, tracing a path of remarkable technological advancements. Introduced by Alan Turing in 1950, the Turing Test challenged the ability to distinguish between responses from a machine and a human during a natural language conversation, setting a foundational criterion for machine intelligence.⁸ In the following decade, significant progress was made with the development of symbolic methods to process natural language, such as the ELIZA chatbot in the 1960s.⁹ In the medical field during the 1970s, expert systems such as Stanford University's MYCIN demonstrated capabilities in reasoning under uncertainty and providing decision rationales, using rule-based AI.¹⁰

The 1990s ushered in the era of ML algorithms with the development of techniques such as support vector machines and random forests. These methods excelled in tasks across various domains by effectively handling pattern recognition and decision-making processes. This period marked a shift from rule-based to data-driven AI, setting the stage for the next big leap. The introduction of deep learning in the 2000s further revolutionized AI capabilities.¹¹ Deep learning models based on neural networks enabled a wide range of complex applications from image recognition to natural language processing. It also led to breakthroughs such as IBM Watson and AlphaGo—systems that famously outperformed humans in Jeopardy and Go, respectively.^{12,13} Although IBM Watson was not eventually successful in healthcare applications, AlphaGo led to the development of AlphaFold, which revolutionized the field of structural biology by solving the protein folding problem.^{14,15}

The emergence of generative AI technologies, especially the LLMs from OpenAI's Generative Pre-trained Transformer (GPT) series and the debut of ChatGPT in November 2022, marked a paradigm shift in the field of AI. Generative AI focuses on creating new content (eg, text, images, and videos) by learning from large existing data, and it differs substantially from earlier AI technologies in several key aspects. First, LLMs use the innovative transformer architecture, which features self-attention mechanisms to evaluate the relevance of each word in a sentence. This enhances the model's grasp of context and linguistic subtleties.¹⁶ Pretrained

on vast text corpora, LLMs utilize billions of parameters to excel across a variety of tasks, such as summarization, translation, question answering, and code generation. These models are often called "foundation models" as the fundamental basis for a wide variety of downstream tasks. Since the development of ChatGPT, numerous LLMs have been created, consistently setting new benchmarks in performance across many tasks.⁷ Additionally, these generative AI technologies offer user-friendly interfaces that allow for direct interaction through natural-language text or speech, significantly reducing the need for structured data preparation typical of earlier ML models such as support vector machines. LLMs can also quickly adapt to new tasks through natural language instructions and in-context learning.

In summary, generative AI technologies such as LLMs represent a fundamental shift in how machines can understand and interact with human language, offering a breadth of capabilities and flexibility that were not possible with earlier AI models. Many applications of LLMs to different data modalities (eg, text, images, and videos) have been investigated across different industrial sectors including healthcare and life science.^{17,18}

Applications of Generative AI in Literature Reviews and Evidence Synthesis

This section identifies applications of generative AI, including the use of foundation models such as LLMs in systematic literature reviews (SLRs) and meta-analyses. At this phase of development, foundation models should be understood as ways to augment human tasks and not as autonomous replacements for humans. There are several aspects of SLRs that can be augmented with foundation models. LLMs can assist in generating a search strategy by proposing Mesh terms and keywords to input in biomedical search engines such as PubMed.¹⁹ Researchers have evaluated the capacities of foundation models to automate abstract and full text screening. When screening full text literature using highly reliable prompts, which are precise inputs guiding the LLM to produce accurate and relevant responses, GPT-4 has been shown to have comparable performance to humans.²⁰ For example, Guo et al²¹ assessed the performance of OpenAI GPT (GPT-3.5 Turbo and GPT-4) in automating screening and found a high level of agreement between the model-generated results and a human reviewer. They also found that the model was able to explain its reasoning for excluding articles and correct its initial decisions.²¹ Robinson et al²² have also evaluated LLMs' capacity to provide exclusion reasoning for abstracts and full texts. Hasan et al²³ compared the agreement between GPT-4 and human reviewers in assessing the risk of bias using the Cochrane Collaboration's risk of bias tool. Their case study demonstrated a fair agreement between the LLM and the human reviewer.²³ Several studies have tested the data extraction capacities of LLMs. Reason et al²⁴ found that GPT-4 >99% accuracy in replicating the data extraction in 4 network meta-analyses. Gartlehner et al²⁵ assessed the performance of CLAUDE-2 in extracting data elements from published studies compared with human extraction and found that CLAUDE-2 achieved an overall accuracy of 96.3%. This study also provides examples of prompt engineering. Schopow et al²⁶ found a high degree of concordance between ChatGPT (GPT-3.5 and 4.0) and human researchers on the data extraction task. Nevertheless, using GPT-4, another study only found moderate performance in assisting with data abstraction.²⁰ In addition to data extraction, the foundation model can also be directed to generate the code to conduct meta-analyses. Reason et al²⁴ found that the LLM used was able to generate error-free code in R for network meta-analyses. However, Qureshi et al¹⁹ found errors in code generated in an earlier GPT model, GPT-3.5. It should be noted that the

presence of errors may be partly attributed to earlier versions of GPT but that other factors may also contribute to this, such as ambiguous or incomplete prompts and the inherent complexity of the coding tasks. Finally, report writing has also been evaluated with researchers finding that the LLM is able to generate reasonable drafts.²⁴

In summary, there is promise in using foundation models to support a range of tasks required in SLRs but the use of LLMs to augment human work in SLRs presents several limitations. These include potential inaccuracies in the outputs, such as errors in classifying abstracts or in extracting data, as well as hallucinations, which are instances in which the output generated is factually incorrect or fabricated. Because LLM-generated citations might be fabricated, manually verifying the terms proposed by the models is required.^{19,27} Search engines can improve their search returns with the integration of AI supported tools and the use of retrieval-augmented generation, a retrieval system to find relevant information to produce more contextually accurate and factual outputs, and which anchors the model's response to verifiable reference.²⁸ Although some studies have demonstrated that LLMs can achieve reasonable accuracy compared with human-performed tasks, this is not universally true, underscoring the necessity for continuous human oversight and validation.²⁷ Other limitations include difficulties in replicating results across different LLMs, which can be due to different user inputs or to the systems themselves. Currently Gen AI tools should be used to augment human tasks not autonomously replace them.

Applications of Generative AI to RWE

Increasingly, RWE is being used by a wide range of stakeholders in the HTA process.²⁹ The potential benefits of using generative AI in RWE include increased efficiency in data processing and analysis and enhanced accuracy and consistency by minimizing human errors and standardizing the evidence generation processes. Unstructured notes in electronic health records, such as radiology reports and physician notes, are important for RWE generation,³⁰ and the use of LLMs can make information extraction from notes easier and more efficient.³¹ A technique known as “few-shot learning,” in which a model is given a few examples in the prompt to guide its response, has been shown to be useful to extract variables from unstructured notes.³² Domain-specific LLMs, such as GatorTron, GatorTronGPT, and Me LLaMA, which are trained using large clinical texts, have been shown to improve the accuracy of outputs.³³⁻³⁵ Additionally, generative AI is increasingly being customized for images and other formats of real-world data to generate evidence.³⁶

There are still limitations when applying LLMs to the area of RWE. These include the generation of inaccuracies and fabrications. For example, one study found that a range of foundation models had <50% accuracy in mapping descriptive text to the correct International Classification of Diseases and current procedural terminology codes.³⁷ Potential for bias and privacy risk will further be discussed in the section on limitations.

Applications of Generative AI to Health Economic Modeling

Generative AI has the potential to support different phases of economic model development, including model conceptualization, parameterization, model implementation, and evaluation and validation of model results.³⁸⁻⁴⁰ Chhatwal et al⁴¹ explored the use of foundation models—GPT-4 and Bing Chat—in the conceptualization of a Markov model for hepatitis C and parameterization of the model. They found significant variability in how these models conceptualized disease progression. The study emphasized the

significance of expert guidance in utilizing LLMs for HEOR model development to ensure accuracy and reliability of the results.

Reason et al⁴² demonstrated the capability of GPT-4 to recreate published 3-state partition survival models for non-small cell lung cancer and renal cell carcinoma. This study provides an in-depth use of prompt engineering, including chain-of-thought prompting, to improve results. Although the model construction was replicated with high accuracy, this approach required human intervention to provide contextual information on top of information specifying the model assumptions, methods, and parameter values. This suggests that, although generative AI can create a model code, significant human expertise is still necessary to guide the model conceptualization and parameterization. In contrast, a study by Ayer et al⁴³ demonstrated the proof-of-concept of fully replicating a published simple health economic model, including extraction of model structure and parameters, code development, and evaluation of results without human involvement. However, further research is needed for fully replicating complex models.

Generative AI could also play an important role in structural uncertainty analysis, a typically resource-intensive process.³⁹ By automating this aspect, generative AI can help in identifying and mitigating potential weaknesses in model structure more efficiently.

Automation of different stages of model development could lead to more efficient model development, reducing time and resource expenditure in the future, but current AI tools will augment and support human experts rather than replacing them entirely. Researchers should be aware of the limitations when using LLMs to in the development of health economic models. These include the potential for inaccuracies, such as erroneous parameters or coding errors, and the potential for fabricated data or parameters.

Limitations of Generative AI in HTA Applications


Although generative AI offers promising applications in areas relevant to HTA, caution should be exercised when evaluating studies using these technologies.⁴ Several frameworks have been developed to help guide the responsible use of AI in healthcare, many of which remain relevant to generative AI.⁴⁴⁻⁴⁶ Here, we briefly review and focus on 3 different categories of limitations of generative AI and foundation models: (1) scientific validity and reliability, (2) bias, equity, and fairness, and (3) regulatory and ethical considerations.

Scientific validity and reliability

As outlined in the introduction, we recommend that Gen AI tools should currently be used to support human activities with humans remaining accountable. Although researchers are responsible for the results' accuracy and reporting, foundation models introduce a new level of complexity that must be carefully managed during analysis and reporting. In the following section, we review some features of foundations models that might affect scientific validity and reliability.

Because LLMs are trained on large corpora of publicly available information from the web, errors may be introduced especially in domain-specific areas, such as health,⁴ requiring careful review of results. A well-noted limitation has been the generation of “hallucinations.”^{27,47} This phenomenon occurs because models generate context based on learned statistical associations between words and sentences, not from an understanding of the knowledge in the corpora. Several strategies might be used to improve the factual correctness of model outputs. First, improving the instructions used in the prompts provided to the models, a process known as prompt engineering, has been shown to improve the

Figure 1. Relationship between AI, Gen AI, FMs, and LLMs.



AI indicates artificial intelligence; FM, foundation model; Gen AI, generative artificial intelligence; GPT, generative pre-trained transformer; LLM, large language model.

accuracy of the results. For example, chain-of-thought prompting is a technique in which a model is guided to generate step-by-step explanation or its reasoning process while arriving at an answer.⁴⁸ Few-shot learning, an in-context learning technology that prompts an LLM with several examples of a specific task, is another approach to improve model performance.⁴⁹ Second, the use of retrieval-augmented generation allows the model to retrieve relevant information from external sources.²⁸ This has been shown reduce the incidence of hallucinations and factually incorrect output.⁵⁰ Third, foundation models can be fine-tuned with domain-specific data and knowledge, thus improving their performance on answering domain-specific questions.^{33,51,52}

In science, findings should also be reproducible, and foundation models add a level of complexity because they are more opaque than traditional statistical tools and learn patterns from data, which can introduce variability in the generation of the model parameters. Model output variation can stem from either user factors (eg, inexperience working with foundation models, suboptimal prompts) or inherent differences in the AI models. Researchers have proposed various approaches to support the reproducibility of studies and the accuracy of results. For example, Hasan et al²³ propose a framework for integrating foundation models into SLRs. Reason et al²⁴ run their models over 20 times to compare the results from different runs. Open sharing of data, code, and results, as well as the adoption of standards for reporting and transparency will likely help improve reproducibility.^{53,54} Finally, transparency in this area is challenging because of the “black-box” nature of LLMs, meaning it is unclear how models generate answers from inputs. Strategies proposed to improve transparency have included requirements for the outputs to cite which part of the data set contributed to the answer and “explainable” AI.⁴ To achieve these goals and address the challenges above, the involvement of all stakeholders, with consideration of human-AI teaming, including human oversight, is critically needed in using these models.⁵⁵

Bias, equity, and fairness

Foundation models can propagate or amplify biases introduced at many steps from model inception to deployment, resulting in systematic differences that may exacerbate inequities and cause harms to individuals and communities.^{56,57} First, systemic bias (or institutional or historical bias) may occur in the use of historical data to train foundation models, in which marginalized groups may be underrepresented in data sets because of barriers to access to resources, perpetuating existing systemic biases in society.⁵⁸ Second, computational and statistical biases may stem from errors that occur when the sample is not representative of the population and the algorithm produces a result that differs from the true underlying estimate.^{59,60} Bias may also occur when specific population groups are excluded from data collection, training, testing, or subsequent analyses.⁵⁶

Several strategies have been proposed to assess and mitigate the risk of bias in foundation models. Surveys on this topic have been published.^{61–63} One review of methods for addressing AI fairness and bias in biomedicine identified distributional and algorithmic approaches for addressing bias.⁶¹ Distributional approaches include strategies such as data augmentation, data perturbation, and data reweighting, such as the generation of synthetic data sets to ensure balanced representation of all demographic groups.⁶⁴ Federated learning may offer some potential to address biases, in which models are trained on multiinstitutional data without the need for data sharing, with the idea that access to more diverse populations in many different health systems may mitigate biases inherent to localized data sets.⁵⁶ Algorithmic approaches change the algorithm itself and include approaches such as adversarial learning and loss-based methods, in which the loss function is adjusted to penalize biased predictions.⁶¹

Regulatory and ethical considerations

Regulatory frameworks focused on generative AI in biomedical research are under development in many countries. However, established data privacy and confidentiality laws, such as the Health Insurance Portability and Accountability Act in the United States and General Data Protection Regulation in the European Union, remain applicable. Gen AI models require a large amount of training data. It is important to avoid using data containing protected health information because absolute deidentification is not attainable, and reidentification risks remain nonzero.^{65,66}

Furthermore, Gen AI models have the ability to memorize the data that they were trained on, posing a risk of reproducing protected health information that they memorized during training.⁶⁷ Using patient-level data or other sensitive, confidential information in commercial LLMs poses risks if data handling and compliance with privacy regulations are unclear because the company providing the models may have access for further model training or other purposes. In contrast, with open-source models, the risk lies in the operator’s control over data security, in which improper handling could lead to leaks. In all cases, clear documentation and data governance must be in place when using public or proprietary datasets for AI training. Enhancements such as synthetic data that replicates real patient data without personal identifiers⁶⁸ and enabling computations on encrypted data⁶⁹ are being explored to augment privacy protections.

The issue of informed consent for AI driven research is becoming increasingly critical and is underscored by recent legal disputes over the use of patient electronic health record data to train LLMs.⁷⁰ Inaccurate information can lead to patient harm, for example, if used to support an HTA submission that erroneously exclude underserved populations. In addition, emergent

BOX 1. Key definitions*.**Generative AI (Gen AI)**

AI systems capable of generating text, images, or other content based on input data, often creating new and original outputs.

- Example: DALL-E generating images from text descriptions.

Foundation Models

Large-scale pretrained models that serve a variety of purposes. These models are trained on broad data at scale and can adapt to a wide range of tasks and domains with or without further fine-tuning. Although LLMs are a subset of foundation models focusing specifically on language, foundation models can encompass other modalities, such as vision and multimodal tasks.

- Examples: GPT-4, Gemini, BERT, CLIP (Contrastive Language-Image Pretraining), and LLama, LLMs

A specific type of foundation model trained on massive text data that can recognize, summarize, translate, predict, and generate text and other content based on knowledge gained from massive data sets.

- Examples: GPT-4, Gemini, BERT, LLama, GatorTron.

AI indicates artificial intelligence; BERT, bidirectional encoder representations from transformers; DALL-E, decoder-only autoregressive language and image synthesis; Gen AI, generative artificial intelligence; GPT, Generative Pre-trained Transformer; LLM, Large Language Model.

*First draft of definitions developed using Chat-GPT4o, May 2024.

capabilities of LLMs may raise significant ethical challenges, particularly when the model's accurate performance cannot be transparently explained. For example, Gichoya et al⁷¹ found that the AI model could identify a patient's self-reported race from imaging alone, across multiple imaging modalities, and often when clinical experts could not. The study was not able to isolate which features the models were using to make these accurate predictions. This raises important ethical issues that will need to be thoughtfully addressed because the recognition by AI models of race in medical images could produce different health outcomes for members of different racial groups.⁷¹ In summary, given the early stages of the development and advances of generative AI, the themes discussed above will continue to be important and the object of warranted attention from many stakeholders within society.

Policy Landscape

Policies regulating generative AI, including foundation models in healthcare, aim to ensure safe, ethical, and effective use of these technologies. The Supplemental Material found at <https://doi.org/10.1016/j.jval.2024.10.3846> includes a table that presents legal and policy documents related to the use of AI, generative AI, and foundation models from governments, regulatory bodies, and multistakeholder groups. The regulations address various aspects, such as risk management, accountability, transparency, and compliance with existing legal standards. The primary objective is to mitigate potential risks associated with AI, such as bias, privacy violations, and safety concerns, while fostering innovation and trust in AI systems.

Different regions adopt varied approaches: the European Union implements a comprehensive, horizontal framework with stringent requirements for high-risk AI,⁷² whereas the United Kingdom, Japan, and the United States favor sector-specific guidelines and flexible regulations.⁷³⁻⁷⁷

BOX 2. Suggestions to HTA agencies and policy makers on the use of LLMs.

- Develop guidance for assessing the appropriate use of LLMs for health technology assessment or endorse existing guidance.
- Provide examples of what constitutes acceptable use of LLMs in health technology assessments, including through use cases and “non-use cases,” ie, areas in which it is not recommended to use LLMs in HTAs.
- Collaborate with other HTA agencies and regulatory agencies to harmonize standards, processes, and reporting requirements for the use of LLMs for health technology assessment.
- Develop guidance on and/or checklists for reporting the use of LLMs, including providing checks on validation, reproducibility, bias. These need not be stand-alone and might be incorporated into sections of existing methodological guidances.
- Assess whether LLM qualification or accreditation schemes, by vendor or data set, would improve the efficiency and quality of submissions to HTA agencies.
- Engage in multistakeholder collaborations to support shared learning around the use of LLMs for HTA, develop best practices, and include the patient perspective.
- Ensure that health equity considerations are appropriately included in the guidance around the use of LLMs in HTA, such as ensuring that populations that have been historically marginalized are represented.
- Invest in training and preparing the HTA workforce and partners (eg, academic collaborators, committee members) in the area of LLMs. Involve relevant experts where possible.

HTA indicates health technology assessment; LLM, Large Language Model.

International organizations have also issued guidance or statements on AI, including the Organization for Economic Co-operation and Development, the United Nations, and United Nations Educational, Scientific, and Cultural Organization.⁷⁸⁻⁸⁰ In the healthcare sector, the World Health Organization has provided guidance to member states on deploying LLMs in healthcare and outlining the risks, such as the overestimating the benefits of using LLMs, propagating system-wide biases, negatively impacting the labor market, and increasing cybersecurity risks.^{81,82}

In the regulatory space, the European Medicines Agency has issued a reflection article on the use of AI/ML in the medicinal product life cycle, including product development, authorization, and postauthorization, emphasizing the need to meet all existing legal requirements, use a risk-based approach, and promote AI trustworthiness.⁸³ The Food and Drug Administration has issued several discussion documents related to AI, specifically on the use of AI/ML-based software as a medical device^{84,85} and more generally on AI use—not specific to generative AI—across the drug development process, including drug discovery, nonclinical phase of drug development, the clinical phase including in clinical trials, postmarketing surveillance, and in manufacturing.⁸⁶⁻⁸⁸

In the field of HTA, we reviewed the websites of several HTA bodies, including the Department of Health and Aged Care in

Australia, the Canadian Agency for Drugs and Technologies in Health, the Haute Autorité de Santé in France, the Institute for Clinical and Economic Review in the United States, the Institut für Qualität und Wirtschaftlichkeit im Gesundheitswesen in Germany, the National Institute for Health and Care Excellence in England, and the Zorginstituut Nederland in The Netherlands. As of the time of writing (October 2024), only National Institute for Health and Care Excellence had issued a position statement on the use of AI in evidence generation.⁸⁹ The ISPOR HTA Roundtable conducted an informal survey with 18 members from North America, Latin America, Australasia, and Europe in April 2024 to gauge the level of internal use of generative AI in HTA agencies (they were not queried about the use of AI in submissions). Several agencies mentioned that they were still in an exploratory phase with generative AI, seeking to better understand how it may support their work and be used by manufacturers in their submissions. Some groups were not using AI at this time and one agency had blocked the use of AI tools such as Chat-GPT at this time. Several agencies are evaluating software offered by private companies in the space of SLRs (5 agencies) and one agency described experimenting with the use of generative AI in economic modeling. Finally, some agencies mentioned internal use of AI tools for translation and for general education purposes. The learnings from these explorations of AI tools will likely inform the methods and processes guidance that HTA agencies will need to develop to guide companies in using generative AI in HTA submissions.⁹⁰

We provide some suggestions for HTA agencies and policy makers for establishing procedures and processes to improve the use of generative AI in HTA submissions based on our experience in the HTA field (Box 2). We recommend establishing clear guidance on the appropriate use of LLMs, including examples of both suitable and unsuitable applications. Standardizing and harmonizing processes across agencies and ensuring transparent reporting will be important. A focus on health equity will enhance the quality and efficiency of submissions and ensure that the appropriate populations are represented in HTA analyses. Additionally, investing in training for the HTA workforce will be essential for the responsible and inclusive use of generative AI technologies.

Multistakeholder groups are also working on frameworks for the responsible and trustworthy use of AI, including, for example, the National Academies of Medicine, the US National Institute of Standards and Technology, the Coalition for Health AI, and the European Commission's HTx H2020 (Next Generation Health Technology Assessment) project.^{44-46,60,90} Finally, we note that scientific and medical journals are in the process of developing their policies with respect to the use of AI tools in publications submitted to their journals.⁹¹ Of note, the new *NEJM AI* journal encourages the use of LLMs in submissions with the assumption that they will enhance the quality of research and support its democratization.⁹²

Conclusion

The field of generative AI provides potentially transformative tools to augment and support the efficient generation of evidence to support HTA under human supervision. In this article, we have provided a brief overview of the development of generative AI and reviewed applications relevant to HTA, including literature review and evidence synthesis, RWE, and health economic modeling. Despite their promise, it is important to acknowledge that these technologies, although rapidly improving, are still nascent and continued careful evaluation in their application to HTA is

required. To ensure the responsible use and implementation of these tools, both developers and users of research must fully understand their limitations, including challenges related to scientific validity and reliability, risks of bias, potential impacts on equity, and critical regulatory and ethical considerations. Because of the nature of these models, there is added complexity in using them and evaluating the validity and reproducibility of their results. It is expected, as with all new technologies, that both user expertise with foundation models and the actual performance of foundation models themselves will improve rapidly in the near future. This is a transformative and exciting time in scientific and medical research, but scientific research is a complex endeavor that affects human health and will continue to require careful validation and human oversight for the foreseeable future.

Glossary

- **AI:** a broad field of computer science that aims to create intelligent machines capable of performing tasks typically requiring human intelligence.
- **Chain-of-thought prompting:** A technique in which the model is guided to reason through a problem step-by-step in its response, by breaking down complex tasks into simpler parts to improve accuracy.
- **Deep Learning:** a subset of ML algorithms that uses multilayered neural networks, called deep neural networks. Those algorithms are the core behind the majority of advanced AI models.
- **Few-Shot Learning:** A method where a model is given a few examples in the prompt to guide its response, utilizing its pretrained knowledge to produce accurate outputs with minimal data.
- **Foundation model:** large-scale pretrained models that serve a variety of purposes. These models are trained on broad data at scale and can adapt to a wide range of tasks and domains with further fine-tuning.
- **Generative AI:** AI systems capable of generating text, images, or other content based on input data, often creating new and original outputs.
- **GPT:** a specific series of LLMs created by OpenAI based on the transformer architecture, which is particularly well suited for generating human-like text.
- **Hallucination:** incorrect output produced by a generative AI model that is not based on the input data or reality. This content is factually incorrect, misleading, or fabricated.
- **LLM:** a specific type of foundation model trained on massive text data that can recognize, summarize, translate, predict, and generate text and other content based on knowledge gained from massive datasets.
- **ML:** a field of study within AI that focuses on developing algorithms that can learn from data without being explicitly programmed.
- **Multimodal AI:** an AI model that simultaneously integrates diverse data formats provided as training and prompt inputs, including images, text, bio-signals, -omics data, and more.
- **Prompt:** the input given to an AI system, consisting of text or parameters that guide the AI to generate text, images, or other outputs in response.
- **Prompt engineering:** creating and adapting prompts (input) to instruct AI models to generate specific output.
- **Retrieval-Augmented Generation:** A method in natural language processing that combines a generative model with a retrieval system to improve response accuracy. The retrieval system finds relevant information, which the generative model uses to produce more contextually accurate and factual outputs.

Author Disclosures

Author disclosure forms can be accessed below in the [Supplemental Material](#) section.

The views expressed are those of the authors and not those of their employing or funding organizations.

Supplemental Material

Supplementary data associated with this article can be found in the online version at <https://doi.org/10.1016/j.jval.2024.10.3846>. found at <https://doi.org/10.1016/j.jval.2024.10.3846>.

Article and Author Information

Accepted for Publication: October 22, 2024

Published Online: December 1, 2024

doi: <https://doi.org/10.1016/j.jval.2024.10.3846>

Author Affiliations: Office of the Director, National Institutes of Health, National Institute of Biomedical Imaging and Bioengineering, Bethesda, MD, USA (Fleurence); Health Outcomes and Biomedical Informatics, College of Medicine, University of Florida, Gainesville, FL, USA (Bian); Biomedical Informatics, Clinical and Translational Science Institute, University of Florida, Gainesville, FL, USA (Bian); Office of Data Science and Research Implementation, University of Florida Health, Gainesville, FL, USA (Bian); Tulane University School of Public Health and Tropical Medicine, New Orleans, LA, USA (Wang); Intelligent Medical Objects, Rosemont, IL, USA (Wang); Department of Biomedical Informatics and Data Science, School of Medicine, Yale University, New Haven, CT, USA (Xu); National Institute for Health and Care Excellence, London, England, UK (Dawoud); Faculty of Pharmacy, Cairo University, Cairo, Egypt (Dawoud); ISPOR-The Professional Society for Health Economics and Outcomes Research, Lawrenceville, NJ, USA (Higashi); Institute for Technology Assessment, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA (Chhatwal); Center for Health Decision Science, Harvard University, Boston, MA, USA (Chhatwal).

Correspondence: Rachael L. Fleurence, PhD, Office of the Director, National Institutes of Health, National Institute of Biomedical Imaging and Bioengineering, Bethesda, MD 20892, USA. Email: rachael.fleurence@nih.gov

Authorship Confirmation: All authors certify that they meet the ICMJE criteria for authorship.

Drs Fleurence, Dawoud, and Chhatwal are editors for *Value in Health* and had no role in the peer-review process of this article.

Funding/Support: Dr Dawoud was partly supported by funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 825162 (Next Generation Health Technology Assessment [HTx] project). The other authors received no financial support for this research.

Role of the Funder/Sponsor: The funder had no role in the design and conduct of the study; collection, management, analysis, and interpretation of the data; preparation, review, or approval of the manuscript; and decision to submit the manuscript for publication.

Acknowledgment: The authors wish to thank the ISPOR Science office for their support as well as Sahar Alam (ISPOR staff) for her excellent program management throughout the project and Kelly Lenahan (ISPOR staff) for conducting a survey of HTA members. ChatGPT-4o was used to generate the first drafts of the abstract, highlights and definitions in [Box 2](#). The authors are responsible for the final content.

REFERENCES

1. What is artificial intelligence? Professor John McCarthy. <http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html>. Accessed May 21, 2024.
2. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. *Nat Med*. 2022;28(1):31–38.
3. Howell MD, Corrado GS, DeSalvo KB. Three epochs of artificial intelligence in health care. *JAMA*. 2024;331(3):242–244.
4. Thirunavukarasu AJ, Ting DSJ, Elangovan K, Gutierrez L, Tan TF, Ting DSW. Large language models in medicine. *Nat Med*. 2023;29(8):1930–1940.
5. Telenti A, Auli M, Hie BL, Maher C, Saria S, Ioannidis JPA. Large language models for science and medicine. *Eur J Clin Invest*. 2024:e14183.
6. Could AI transform life in developing countries? *The Economist*. <https://www.economist.com/briefing/2024/01/25/could-ai-transform-life-in-developing-countries>. Accessed November 22, 2024.
7. Zhao WX, Zhou K, Li J, et al. A survey of large language models. arXiv. <https://arxiv.org/pdf/2303.18223>. Accessed November 22, 2024.
8. Turing AM. *Computing Machinery and Intelligence*. Berlin, Germany: Springer; 2009.
9. Weizenbaum J. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun ACM*. 1966;9(1):36–45.
10. Shortliffe E. MYCIN: a rule-based computer program for advising physicians regarding Antimicrobial Therapy selection 1976. 1976;01/01, 739-739 https://www.researchgate.net/publication/234816828_MYCIN_A_Rule-Based_Computer_Program_For_Advising_Physicians_Regarding_Antimicrobial_Therapy_Selection. Accessed November 22, 2024.
11. Hinton GE, Osindero S, Teh YW. A fast learning algorithm for deep belief nets. *Neural Comput*. 2006;18(7):1527–1554.
12. Ferrucci DA. Introduction to “this is Watson”. *IBM J Res Dev*. 2012;56(3.4):1.1–1.15.
13. Silver D, Huang A, Maddison CJ, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*. 2016;529(7587):484–489.
14. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583–589.
15. Remarkable progress has been made in understanding the folding of proteins. *The Economist*. <https://www.economist.com/leaders/2021/07/31/remarkable-progress-has-been-made-in-understanding-the-folding-of-proteins>. Accessed November 22, 2024.
16. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. *Adv Neural Inf Process Syst*. 2017;30:1–11.
17. Moor M, Banerjee O, Abad ZSH, et al. Foundation models for generalist medical artificial intelligence. *Nature*. 2023;616(7956):259–265.
18. AIs will make health care safer and better. *The Economist*. <https://www.economist.com/technology-quarterly/2024/03/27/ais-will-make-health-care-safer-and-better>. Accessed November 22, 2024.
19. Qureshi R, Shaughnessy D, Gill KAR, Robinson KA, Li T, Agai E. Are ChatGPT and large language models “the answer” to bringing us closer to systematic review automation? *Syst Rev*. 2023;12(1):72.
20. Khraisha Q, Put S, Kappenberg J, Warratich A, Hadfield K. Can large language models replace humans in systematic reviews? Evaluating GPT-4's efficacy in screening and extracting data from peer-reviewed and grey literature in multiple languages. *Res Synth Methods*. 2024;15(4):616–626.
21. Guo E, Gupta M, Deng J, Park YJ, Paget M, Naugler C. Automated paper screening for clinical reviews using large language models: data analysis study. *J Med Internet Res*. 2024;26:e48996.
22. Robinson A, Thorne W, Wu BP, et al. Bio-SIEVE: exploring instruction tuning large language models for systematic review automation. arXiv. <https://arxiv.org/pdf/2308.06610>. Accessed November 22, 2024.
23. Hasan B, Saadi S, Rajjoub NS, et al. Integrating large language models in systematic reviews: a framework and case study using ROBINS-I for risk of bias assessment. *BMJ Evid Based Med*. Published online February 21, 2024. <https://doi.org/10.1136/bmjebm-2023-112597>.
24. Reason T, Benbow E, Langham J, Gimblett A, Klijn SL, Malcolm B. Artificial intelligence to automate network meta-analyses: four case studies to evaluate the potential application of large language models. *Pharmacoecon Open*. 2024;8(2):205–220.
25. Gartlehner G, Kahwati L, Hilscher R, et al. Data extraction for evidence synthesis using a large language model: a proof-of-concept study. *Res Synth Methods*. 2024;15(4):576–589.
26. Schopow N, Osterhoff G, Baur D. Applications of the natural language processing tool ChatGPT in clinical practice: comparative study and augmented systematic review. *JMIR Med Inform*. 2023;11:e48933.
27. Jin Q, Leaman R, Lu Z. Retrieve, summarize, and verify: how will ChatGPT affect information seeking from the medical literature? *J Am Soc Nephrol*. 2023;34(8):1302–1304.
28. Wei CH, Allot A, Lai PT, et al. PubTator 3.0: an AI-powered literature resource for unlocking biomedical knowledge. *Nucleic Acids Res*. 2024;52(W1):W540–W546.
29. Fleurence R, Kent S, Adamson B, et al. Assessing real-world data from electronic health records for health technology assessment - the SUITABILITY

- checklist: a good practices report from an ISPOR task force. *Value Health*. 2024;27(6):692–701.
30. Keloth VK, Banda JM, Gurley M, et al. Representing and utilizing clinical textual data for real world studies: an OHDSI approach. *J Biomed Inform*. 2023;142:104343.
 31. Hu Y, Chen Q, Du J, et al. Improving large language models for clinical named entity recognition via prompt engineering. *J Am Med Inform Assoc*. 2024;31(9):1812–1820.
 32. Agrawal M, Heggelmann S, Lang H, Kim Y, Sontag D. Large language models are few-shot clinical information extractors. arXiv. <https://arxiv.org/pdf/2205.12689>. Accessed November 22, 2024.
 33. Yang X, Chen A, PourNejatian N, et al. A large language model for electronic health records. *NPJ Digit Med*. 2022;5(1):194.
 34. Peng C, Yang X, Chen A, et al. A study of generative large language model for medical research and healthcare. *NPJ Digit Med*. 2023;6(1):210.
 35. Xie Q, Chen Q, Chen A, et al. Me-LLaMa: foundation large language models for medical applications. ArXiv. <https://arxiv.org/pdf/2402.12749>. Accessed November 22, 2024.
 36. Xu H, Usuyama N, Bagga J, et al. A whole-slide foundation model for digital pathology from real-world data. *Nature*. 2024;630(8015):181–188.
 37. Soroush A, Glicksberg BS, Zimlichman E, et al. Large language models are poor medical coders – benchmarking of medical code querying. *NEJM AI*. 2024;1(5):Aldbp2300040.
 38. Caro JJ, Briggs AH, Siebert U, Kuntz KM, ISPOR-SMDM Modeling Good Research Practices Task Force. Modeling good research practices—overview: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force–1. *Value Health*. 2012;15(6):796–803.
 39. Briggs AH, Weinstein MC, Fenwick EA, et al. Model parameter estimation and uncertainty: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force–6. *Value Health*. 2012;15(6):835–842.
 40. Eddy DM, Hollingworth W, Caro JJ, et al. Model transparency and validation: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force–7. *Value Health*. 2012;15(6):843–850.
 41. Chhatwal J, Yildirim IF, Balta D, et al. Can large language models generate conceptual health economic models? ISPOR. <https://www.ispor.org/heor-resources/presentations-database/presentation/intl2024-3898/139128>. Accessed November 22, 2024.
 42. Reason T, Rawlinson W, Langham J, Gimblett A, Malcolm B, Klijn S. Artificial intelligence to automate health economic modelling: a case study to evaluate the potential application of large language models. *Pharmacoecon Open*. 2024;8(2):191–203.
 43. Ayer T, Samur S, Yildirim IF, Bayraktar E, Ermis T, Chhatwal J. *Fully replicating published health economic models using generative AI*. Boston, MA: presented at: Annual Meeting of the Society for Medical Decision Making; 2024.
 44. Adams L, Fontaine E, Lin S, Crowell T, Chung VCH, Gonzalez A. Artificial intelligence in health, health care, and biomedical science: an AI code of conduct principles and commitments discussion draft. NAM Perspectives. https://nam.edu/wp-content/uploads/2024/04/Artificial-Intelligence-in-Health-Health-Care-and-Biomedical-Science_final_4.8.24V2.pdf. Accessed November 22, 2024.
 45. Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>. Accessed November 22, 2024.
 46. Blueprint for trustworthy AI implementation guidance and assurance for healthcare. Coalition for Health AI. https://www.coalitionforhealthai.org/papers/blueprint-for-trustworthy-ai_V1.0.pdf. Accessed November 22, 2024.
 47. Alkaiissi H, McFarlane SI. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus*. 2023;15(2):e35179.
 48. Lin Z. How to write effective prompts for large language models. *Nat Hum Behav*. 2024;8(4):611–615.
 49. Wang S, Scells H, Zhuang S, Potthast M, Koopman B, Zuccon G. Zero-shot generative large language models for systematic review screening automation. arXiv. <https://arxiv.org/pdf/2401.06320>. Accessed November 22, 2024.
 50. Yu H, Gan A, Zhang K, Tong S, Liu Q, Liu Z. Evaluation of retrieval-augmented generation: a survey. arXiv. <https://arxiv.org/pdf/2405.07437>. Accessed November 22, 2024.
 51. Saab K, Tu T, Weng WH, et al. Capabilities of Gemini models in medicine. ArXiv. <https://arxiv.org/pdf/2404.18416>. Accessed November 22, 2024.
 52. Yang L, Xu S, Sellergren A, et al. Advancing multimodal medical capabilities of Gemini. ArXiv. <https://arxiv.org/pdf/2405.03162>. Accessed November 22, 2024.
 53. Beam AL, Manrai AK, Ghassemi M. Challenges to the reproducibility of machine learning models in health care. *JAMA*. 2020;323(4):305–306.
 54. Kapoor S, Cantrell EM, Peng K, et al. REFORMS: consensus-based recommendations for machine-learning-based science. *Sci Adv*. 2024;10(18):eadk3452.
 55. National Academies of Sciences Engineering, and Medicine. *Human-AI Teaming: State-of-the-Art and Research Needs*. Washington, DC: The National Academies Press; 2022.
 56. Drukker K, Chen W, Gichoya J, et al. Toward fairness in artificial intelligence for medical image analysis: identification and mitigation of potential biases in the roadmap from data collection to model deployment. *J Med Imaging (Bellingham)*. 2023;10(6):061104.
 57. Gervasi S, Chen I, Smith-McLallen A, et al. The potential for bias in machine learning and opportunities for health insurers to address it. *Health Aff (Millwood)*. 2022;41(2):212–218.
 58. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447–453.
 59. Parikh RB, Teeple S, Navathe AS. Addressing bias in artificial intelligence in health care. *JAMA*. 2019;322(24):2377–2378.
 60. Towards a standard for identifying and managing bias in artificial intelligence. National Institute of Standards and Technology. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>. Accessed November 22, 2024.
 61. Yang Y, Lin M, Zhao H, Peng Y, Huang F, Lu Z. A survey of recent methods for addressing AI fairness and bias in biomedicine. *J Biomed Inform*. 2024;154:104646.
 62. Xu J, Xiao Y, Wang WH, et al. Algorithmic fairness in computational medicine. *EBioMedicine*. 2022;84:104250.
 63. Huang Y, Guo J, Chen WH, et al. A scoping review of fair machine learning techniques when using real-world data. *J Biomed Inform*. 2024;151:104622.
 64. Juwara L, El-Hussuna A, El Emam K. An evaluation of synthetic data augmentation for mitigating covariate bias in health data. *Patterns (N Y)*. 2024;5(4):100946.
 65. Benitez K, Malin B. Evaluating re-identification risks with respect to the HIPAA privacy rule. *J Am Med Inform Assoc*. 2010;17(2):169–177.
 66. Simon GE, Shortreed SM, Coley RY, et al. Assessing and minimizing re-identification risk in research data derived from health care records. *EGEMS (Wash DC)*. 2019;7(1):6.
 67. Li H, Moon JT, Purkayastha S, Celi LA, Trivedi H, Gichoya JW. Ethics of large language models in medicine and medical research. *Lancet Digit Health*. 2023;5(6):e333–e335.
 68. Mosquera L, El Emam K, Ding L, et al. A method for generating synthetic longitudinal health data. *BMC Med Res Methodol*. 2023;23(1):67.
 69. Brannvall R, Forsgren H, Linge H, HEIDA: software examples for rapid introduction of homomorphic encryption for privacy preservation of health data. *Stud Health Technol Inform*. 2023;302:267–271.
 70. Duffourc MN, Gerke S. Health care AI and patient Privacy-Dinerstein v Google. *JAMA*. 2024;331(11):909–910.
 71. Gichoya JW, Banerjee I, Bhimireddy AR, et al. AI recognition of patient race in medical imaging: a modelling study. *Lancet Digit Health*. 2022;4(6):e406–e414.
 72. EU artificial intelligence act. European Union. <https://artificialintelligenceact.eu/the-act/>. Accessed May 22, 2024.
 73. A pro-innovation approach to AI regulation: government response. UK Government. <https://www.gov.uk/government/consultations/ai-regulation-a-pro-innovation-approach-policy-proposals/outcome/a-pro-innovation-approach-to-ai-regulation-government-response>. Accessed November 22, 2024.
 74. A pro-innovation approach to AI regulation. UK Government. <https://assets.publishing.service.gov.uk/media/64cb71a547915a00142a91c4/a-pro-innovation-approach-to-ai-regulation-amended-web-ready.pdf>. Accessed May 22, 2024.
 75. Habuka H. Japan's approach to AI regulation and its impact on the 2023 G7 presidency. G7 2023 Hiroshima Summit. <https://shorturl.at/MQvju>. Accessed November 22, 2024.
 76. The White House. Blueprint for an AI Bill of Rights. <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>. Accessed November 22, 2024.
 77. FACT SHEET. President Biden issues executive order on safe, secure, and trustworthy artificial intelligence. The White House. <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>. Accessed May 22, 2024.
 78. OECD AI Principles overview. Organization for Economic Co-operation and Development (OECD). <https://oecd.ai/en/ai-principles>. Accessed May 22, 2024.
 79. Ethics of artificial intelligence. United Nations Educational, Scientific and Cultural Organization. <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>. Accessed May 22, 2024.
 80. Interim report: governing AI for humanity. United Nations. Accessed November 22, 2024. https://www.un.org/sites/un2.un.org/files/un_ai_advisory_body_governing_ai_for_humanity_interim_report.pdf.
 81. Regulatory considerations on artificial intelligence for health. World Health Organization. <https://iris.who.int/bitstream/handle/10665/373421/9789240078871-eng.pdf?sequence=1>. Accessed November 22, 2024.
 82. Ethics and governance of artificial intelligence for health. Guidance on large multi-modal models/ World Health Organization. <https://iris.who.int/bitstream/handle/10665/375579/9789240084759-eng.pdf?sequence=1>. Accessed May 22, 2024.
 83. Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle. European Medicines Agency. <https://www.ema.europa.eu/en/use-artificial-intelligence-ai-medicinal-product-lifecycle>. Accessed May 22, 2024.
 84. Artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD) action plan. US Food and Drug Administration. <https://www.fda.gov/media/145022/download>. Accessed November 22, 2024.
 85. Predetermined change control plans for machine learning-enabled medical devices: guiding principles. US food and Drug Administration. <https://www.fda.gov/media/173206/download?attachment>. Accessed May 22, 2024.

86. Artificial intelligence and medical products: how CBER, CDER, CDRH, and OCP are working together. US Food and Drug Administration. <https://www.fda.gov/media/177030/download?attachment>. Accessed November 22, 2024.
87. Using artificial intelligence and machine learning in the development of drug and biological products. US Food and Drug Administration. Accessed November 22, 2024. <https://www.fda.gov/media/167973/download>.
88. Artificial intelligence in drug manufacturing. US Food and Drug Administration. <https://www.fda.gov/media/165743/download>. Accessed May 29, 2024.
89. Use of AI in evidence generation: NICE position statement. National Institute for Health and Care Excellence. <https://www.nice.org.uk/about/what-we-do/our-research-work/use-of-ai-in-evidence-generation-nice-position-statement>. Accessed September 20, 2024.
90. Zemplyeni A, Tachkov K, Balkanyi L, et al. Recommendations to overcome barriers to the use of artificial intelligence-driven evidence in health technology assessment. *Front Public Health*. 2023;11:1088121.
91. Flanagan A, Pirracchio R, Khera R, Berkwits M, Hswen Y, Bibbins-Domingo K. Reporting use of AI in research and scholarly publication-JAMA network guidance. *JAMA*. 2024;331(13):1096–1098.
92. Koller D, Beam A, Manrai A, et al. Why we support and encourage the use of large language models in NEJM AI submissions. *NEJM AI*. 2024;1(1):A1e2300128.