



AI That Listens, Documents, and Engages

Shashikala Valiki

Independent Researcher

shashikala.valiki.researcher@gmail.com

Abstract

Generative AI systems, capable of producing coherent, contextually relevant text, images, and other media from prompt inputs, have become increasingly accessible. The potential for Generative AI to improve patient care and clinician efficiency has generated considerable interest in the healthcare sector, focusing on the enhancement of clinical documentation and patient engagement processes. The use of Generative AI in these domains is discussed with a focus on the underlying technology, implementation considerations for healthcare organizations, and case studies demonstrating the effectiveness of Generative AI in real-world deployments.

Automated generation of clinical notes based on free-text summaries, unstructured summaries of patient examinations and assessments, or conversational inputs is explored, along with the code-based structuring of free-text notes and the application of standardization templates to ensure compliance. The generation of patient education materials appropriate for health literacy levels and cultural backgrounds, the scheduling of appointments, and the triaging of patient queries using Generative AI are also covered. Ethical considerations—especially with respect to data governance and the potential for biased, adversarial, or inaccurate output—are flagged throughout, along with the importance of establishing and maintaining high-quality workflows for the use of Generative AI services.

Keywords: Generative Artificial Intelligence In Healthcare, Clinical Documentation Automation, Patient Engagement Systems, Automated Clinical Note Generation, Unstructured Medical Text Processing, Conversational AI In Healthcare, Clinical Workflow Efficiency, Health Literacy-Aware Content Generation, Patient Education Automation, Appointment Scheduling Systems, Intelligent Patient Triage, Healthcare AI Implementation, Clinical Standardization Templates, Data Governance In Healthcare AI, Ethical AI In Medicine, Bias And Risk Management, Adversarial AI Concerns, Compliance In Clinical Documentation, Real-World Healthcare AI Deployments, Human-Centered Clinical AI Systems.

1. Introduction

Generative AI technologies have become the focus of considerable attention in both the public and healthcare communities. In the healthcare sector, the use of Generative AI solutions is broadly anticipated to impact clinical documentation and patient engagement processes. Clinical documentation and patient engagement technologies are the focus of current investigations by various US and international federal government agencies, healthcare consulting firms, and select health organizations. The potential to support these areas by enabling quicker and more accurate clinical note generation, improving coding capture of clinical visits, enhancing patient education and

communication materials, improving appointment scheduling, and triaging patient inquiries is being assessed. A number of questions surrounding these applications of Generative AI are emerging for investigation. Is Generative AI sufficiently reliable in its output to replace human interaction, or is it best suited as a decision support tool? What areas of clinical documentation and patient engagement may benefit from the deployment of Generative AI? Moreover, what strategies should healthcare organizations implement when considering its adoption? Answering these questions, especially in areas clinical documentation and patient engagement, is important for both US and global healthcare systems.



1.1. Overview of the Role of Generative AI in Transforming Healthcare

The potential and promise of generative artificial intelligence (gen AI) are beginning to be realized in healthcare. Specific implementations are appearing in clinical documentation workflow—including note generation, coding, quality assurance, and regulatory compliance—as well as in patient-facing applications such as education, communication, appointment scheduling, and triage.

The impact in healthcare mirrors rapid developments seen across various sectors. Gen AI is defined as a branch of artificial intelligence that is capable of generating text, images, and other information in response to user prompts, making it the first thing that experts say must be considered when evaluating the future of AI. Applications within this area utilize machine learning models that have been trained on massive amounts of data to enable the generation of human-like output. The most obvious healthcare-related implementations involve large language models (LLMs)—AI architectures trained on vast stores of human knowledge that ingests, processes, and creates responses to voluminous natural language queries and prompts—yet the technology is by no means limited to text input and output. Among other aspects of LLMs, ready access to the internet, retrieval-augmented generation (RAG), and natural language understanding are beginning to be employed to create and fulfill patient and clinician needs across all settings within healthcare.




Fig 1: Generative AI in the Clinical Ecosystem: Harnessing LLMs and Retrieval-Augmented

Generation for Integrated Documentation and Patient-Centric Care

2. Foundations of Generative AI in Healthcare

Generative AI relies on several technical underpinnings – large language models (LLMs), retrieval-augmented generation (RAG), and natural language understanding (NLU) – and relies on key concepts such as tokenization, prompting, evaluation metrics, and bias remediation. Understanding how these components work together elucidates the capabilities, limitations, and meaningful application of Generative AI to clinical use cases. Served by LLMs, Generative AI is capable of producing human-quality text for a variety of basic writing tasks. LLMs are trained on natural conversations, coding languages, and textual corpora that span multiple disciplines. While this provides a strong foundation for a wide variety of AI-driven products and services, such capabilities still require careful prompting if reliable results are to be achieved. Verifiably accurate outputs are better produced via merging LLMs with other forms of AI such as NLU, RAG, and specialized pipelines that verify accuracy. Prompts can influence an AI's response by shaping its role in and the constraints of a task, but the format and content of the response itself, and externally guided or retrieved context.


2.1. Core Technologies and Concepts

The foundational technologies of generative artificial intelligence (AI) in healthcare are large language models (LLMs) and retrieval-augmented generation (RAG). A large language model is a neural network trained to predict the next token of a sequence by processing vast data sets containing a mix of unlabelled text, such as books, web pages, and news articles, with no regard for downstream tasks. RAG refers to the coupling of a neural retrieval mechanism, such as one based on embeddings, with a neural generation mechanism, such as a large language model. RAG combines the advantages of both approaches by efficiently adding current and task-specific knowledge to an LLM without incurring the enormous storage and



querying costs associated with storing and manipulating all possible key answers during training.

Three key natural language processing concepts are especially relevant for working with LLMs. Tokenization is the process of splitting text into smaller components known as tokens. A token can correspond to everything from individual characters to whole words (or more) and is usually chosen according to some rate of exchange between the size of the model and the efficiency of tokenization. Prompting is the process of constructing the input to a neural-network generative model in a way that maximises the quality of any resulting output. A five-part structure has been shown to generalise across different tasks and systems: task specification; task elements; examples; proposed solution; and constraints. Evaluation metrics assess the quality of a model's output, typically by quantifying its similarity to trusted ground-truth outputs. Common metrics include BLEU, ROUGE, and BERTScore, all of which measure different aspects of textual similarity. Bias mitigation in LLMs evaluates and reduces bias that emerges in model outputs without affecting the model's performance on tasks not involving sensitive attributes. Biases can be divided into four categories: societal stereotypes; topical preferences; minority under-representation; and probabilistic associations.



Equation 1: Large Language Models (LLMs): Next-token prediction

1.1 Tokenization (concept → math)

A text sequence is converted into tokens:

$$x = (x_1, x_2, \dots, x_T)$$

where each x_t is a token id from a vocabulary V (size $|V|$).

LLMs learn the **conditional probability**:

$$P(x_1, \dots, x_T) = \prod_{t=1}^T P(x_t | x_{<t})$$

3. Generative AI for Clinical Documentation

Generative AI enables the automated generation and structuring of clinical documentation. Natural inputs, such as conversational transcripts and unaudited notes, are actively transformed into accurate, standard-compliant, and regulatory-aligned outputs, while passive sources, including external data, validated prompts, specific templates, and historical patterns, facilitate the safe generation of auto-coded, auditable, and risk-mitigated clinical summaries. Current applications offer seamless integration with electronic health record (EHR) systems. Automated Documentation Generation

Natural inputs, whether conversational transcripts from AI-supported chatbots or unaudited notes from auto-scribing tools, are actively transformed into accurate, standard-compliant, and regulatory-aligned EHR entries. Passive sources actively augment generative processes and proactively ensure safety. External data—such as retrieval-augmented generation, LangChain sequences, user dialogue history, domain-specific patterns, and browse-ready plugins—implement recommended approaches while facilitating the safe utilization of auto-coded, auditable, and risk-mitigated clinical summaries shaped according to clinician-determined templates. Integrative points with EHRs include docs scribes, coding assistants, and cloud-based applications, as well as the offshore support team of every culturally, linguistically, and legally sensitive organization.



3.1. Automated Note Generation and Structuring

Generative AI supports the automated generation and structuring of clinical notes in augmenting clinical documentation. Note generation automatically creates complete documents from patient conversational transcripts, dictations, or other unstructured inputs (and sometimes from structured data), whereas note structuring organizes these documents into standard fields (e.g., assessment, plan) for entry into the electronic health record (EHR). Support for both services is still in early stages with largely experimental offerings, requiring appropriate clinical supervision and oversight. Use cases, nevertheless, continue to evolve, shaping the maturity of these interventions.

Although the core language model can effectively automate note generation and structuring, separate retriever–translator architectures benefit the outcome quality; these systems complement generation with external retrieval sources, such as internal knowledge bases, as well as integration with electronic health record (EHR) systems. The process begins with a synchronous or asynchronous conversation involving a patient and clinician and may include unstructured note-taking by the clinician. The generated notes may take the form of completed documents, ready for addition to the EHR, or a structured output (e.g., an assessment followed by a plan) that aligns with EHR fields. Key inputs include patient-clinician dialogue, supporting clinical data (e.g., pertinent lab findings), and patient records, tests, and surveys made available through the EHR. Notes or note components also can be generated or structured from clinician dictation or other transcripts. Standard EHR format compatibility facilitates direct addition into downstream systems.





Fig 2: Automating the Clinical Narrative: A Retriever–Translator Framework for Generative Note Structuring and EHR Integration

3.2. Coding, Compliance, and Quality Assurance

Automated generation of clinical notes has clear implications for coding, documentation integrity, and fraud detection. Accurate and standardized clinical notes support the assignment of diagnostic and procedural codes that are cost- and resource-utilization–related. Language that meets requirements for risk-bearing healthcare organizations and third-party payers can help ensure clinical notes sustain regulatory scrutiny and stand up to audits for proper assignment of risk-adjusted payments to healthcare providers.

For large language models (LLMs) and related technologies to perform in this area, engagement and quality control processes must be in place. These consider that AI models can fail to produce human-level output quality in their default state. Quality assurance standards assess the quality of content produced using LLMs, and multiple strategies, such as prompting, decoding, and post-processing, can help ensure that finished products meet acceptable standards. In the clinical context, clinicians and other healthcare providers can implement closed-loop processes to enhance and refine model-generated content.



4. Generative AI for Patient Engagement

Patient education and communication are prime candidates for generative AI applications. During an appointment, patients typically retain only 30% to 50% of the information conveyed by physicians. Moreover, patients report that their providers often use medical jargon, leaving them confused despite attempts to ask clarifying questions. Linguistic, cultural, literacy, and socioeconomic barriers hinder patient understanding. Consequently, a sizable percentage of patients utilize the Internet as their primary information source and arrive at consultations with misguided beliefs.

Generative AI models can summarize information about a specific diagnosis, procedure, or medication, which clinicians can review and adjust prior to sharing. Because larger models can tailor output to reflect patients' health literacy level and generate content in multiple languages, they can help ensure that educational material is linguistically and culturally appropriate. Additionally, by relying on its ever-expanding knowledge base, generative AI can expedite the informed consent process by generating descriptions of complex procedures. Inputting brief details about a planned intervention enables the generation of a thorough description in terms the patient can readily understand.

Appointment scheduling and triage represent points in the healthcare experience where generative AI can be integrated to streamline workflows. Considerable research and development have driven the creation of chatbots that assist with appointment requests. Typically, these chatbots employ decision tree-based logic that leads the user through a series of progressively specific questions to help route requests to the correct departments or specialties, identify scheduling availability, and assess the urgency of the appointment. Generative AI takes such systems a step further by enabling natural language queries, alleviating challenges patients face when trying to determine which health concerns warrant medical attention, and flagging potentially urgent cases for immediate review by a clinician. Working in concert with safety measures—such as links to a symptom checker or escalation of specific phrases or words—these capabilities can provide more efficient patient support while safeguarding against harm.

Equation 2: Softmax: converting logits to probabilities (step-by-step)

At a time step t , the model outputs a vector of logits:

$$z_t \in \mathbb{R}^{|V|}$$

Step 1: exponentiate logits

$$\tilde{p}_{t,i} = e^{z_{t,i}}$$

Step 2: normalize so probabilities sum to 1

$$p_{t,i} = \frac{e^{z_{t,i}}}{\sum_{j=1}^{|V|} e^{z_{t,j}}}$$

So:

$$P(x_t = i \mid x_{<t}) = \text{softmax}(z_t)_i$$


4.1. Patient Education and Communication

Generative AI capabilities can enhance patient education materials and other communications by tailoring content to patients' unique circumstances, needs, and preferences. The technology can simplify complex medical terms into easy-to-understand language, adjust cultural references and



health literacy levels to reflect patient backgrounds, and provide multilingual support. These enhancements can facilitate better patient understanding of care and treatment and promote informed consent for surgery, anesthesia, and other procedures.

Generative AI can also facilitate appointment scheduling and pre-visit communications, routing requests to the appropriate healthcare team members while assessing the urgency and safety of requests. The technology can help scale these efforts in support of patient-centered models of access to care. For example, some healthcare systems have tested AI applications for triaging patients who use chat features to relay potential symptoms of an urgent medical issue (e.g., stroke). In those cases, the chatbots recommended that patients seek immediate medical attention at an emergency department instead of scheduling an appointment with a primary care physician. Addressing an objection raised about the safety of such systems, researchers noted that built-in safety rules mitigated the risks inherent in such tasks.



4.2. Appointment Scheduling and Triage

Appointment scheduling and triage systems are essential for managing the volume, frequency, and urgency of patient interactions. Generative AI offers substantial improvements in these areas. Clinical workloads and workflows can be better matched, avoiding abandonment when capacity is exceeded, minimizing unnecessary interactions when capacity is underutilized, and directing patients to the appropriate resource in a timely manner.

Moreover, integration with electronic health record (EHR) systems enables bidirectional clinical contextualization of patient communications.

Generative AI can enhance clinical appointment scheduling by interacting directly with patients and checking availability in real time. Similar to appointment booking in the travel industry, a natural language interface provides convenience and familiarity. Conversational agents can also perform symptom checking and triage. These systems detect the presenting problem, assess the required urgency, and route the request to the appropriate resource (e.g., primary care, specialty care, nursing, or behavioral health). Together with expanded access (e.g., extended hours and virtual visits), advanced triage systems can significantly improve patient outcomes for urgent medical needs, especially among populations that may not otherwise seek early treatment. Recognizing the risk of acute or serious illness can be difficult, particularly in children. Natural language processing applied to online health-related media offers opportunities for automated detection of emerging safety hazards. However, these enhancements must be supported by careful clinical governance.

5. Implementation Strategies for Healthcare Organizations

Implementing generative AI solutions in healthcare organizations requires careful evaluation of potential vendor systems and products, as well as a focus on change management and the readiness of users to adapt to automated processes. Considerations include assessing clinical and operational needs, ensuring interoperability of any selected system, and establishing data governance and security standards.

Vendor offerings must be evaluated against multiple criteria. A focus on the most relevant requirements for a given organization can facilitate a more efficient selection process. Interoperability with other vendor systems and key clinical services is crucial. Similarly, established and enforceable data governance policies and security controls must be in place to ensure data privacy and integrity, detect possible misuse, and manage vendor-related data risks. Change management strategies aimed at optimizing



clinician acceptance of generative systems should also be developed. Supply-chain networks need to ensure vendors meet audit and testing requirements, provide transparent data usage standards, and offer reliable models and high-quality outputs.




Fig 3: Scaling Generative Healthcare: A Strategic Framework for Vendor Governance, Change Management, and Multilingual Patient Engagement

Generative AI opens exciting new avenues for patient engagement, education, and communications. Solutions can offer personalized information about diagnoses, medications, and healing processes in user-friendly, comprehensible formats. Translation services can provide non-English-speaking patients—especially valuable in regions where multiple languages are spoken—with clear instructions, educational material, and informed-consent forms in their native languages. Enhancing formal documentation with official translation and approval workflows retains data integrity and supports compliance. Systems can direct patients to the correct healthcare services, alert providers to new bookings, and provide initial assessments of urgency. Special attention to safeguarding against possible misdirections can help mitigate risk.

Prompt part	Symbol	Purpose	Example (clinical note)
Task specification	<i>T</i>	Define objective	“Generate SOAP note”
Task elements	<i>E</i>	Provide inputs	“Transcript + vitals”
Examples	<i>X</i>	Demonstrate format	Sample SOAP output
Proposed solution	<i>S</i>	Provide draft/plan	“Summarize then structure”
Constraints	<i>C</i>	Safety/compliance	“No hallucinations; cite meds”

5.1. Vendor Assessment and Interoperability

For healthcare organizations, implementing generative AI capabilities requires assessing vendors, ensuring interoperability, complying with data governance and security policies, and managing vendor risk. Organizations should evaluate vendors' claims with proof-of-concept testing and consideration of clinical merit. Generative AI is most likely to enhance existing functionality, and thus, organizations should consider incorporating generative AI in combination with large language model capabilities for end-users via an artificial intelligence support routing layer. Support for retrieval-augmented generation should be prioritized, enabling organizations to supplement large language models with EHR-specific context and knowledge.

Healthcare organizations will recognize established data governance, security, and vendor risk management processes as further requirements in the implementation of generative AI capabilities. To mitigate the unintentional exposure of private health information, vendor risk management teams should pay particular attention to data handling. Adherence to the explicit permission- and-opt-out relationships of the Health Insurance Portability and



Accountability Act will be required to safeguard data at all procurement stages. Openness to risk-sharing is recommended, including fronting the costs for generative AI vendors to build an earlier detection system for verity-transparency issues (the generation of false yet convincing content).

5.2. Change Management and Clinician Adoption


Successful Generative AI implementation requires effective management of the people side of change. Reported failure rates for such initiatives are disturbingly high, with estimates of 70 percent or higher for IT-enabled transformations. Healthcare organizations are not exempt. Some organizations have adopted change accelerators known to mitigate common reasons why change initiatives falter. Key to success is a focus on change management specific to the involved clinical staff, as clinician adoption is essential for effective Generative AI usage. Leaders of Generative AI efforts often view these initiatives strictly as an IT project, even though meaningful impact programs beyond cost reduction are fundamentally about driving behavior change. The need for clinicians to adapt to new workflows using unfamiliar tools and detailed clinical content generated by an AI engine in the background makes these projects as much about change management as about information technology.

Adoption and utilization of Generative AI applications for clinical documentation can be viewed through the lens of the 5Cs. Supported by structured training materials, a strategic pilot approach can drive clinician behavior change. By integrating Generative AI implementations into routine meetings that help distribute news and shape perceptions, stakeholders can influence the 5Cs: clarity, competence, confidence, commitment, and collaboration. Auditing, monitoring, and reporting against adoption and utilization goals can be planned at the outset within the overall evaluation plan. Usability testing remains important for quality assurance and can be done in a staged or accelerated manner.

6. Case Studies and Practical Applications

The implementation of generative AI capabilities in health systems is at an early stage, but initial commercial offerings are becoming available and early work is being reported. GenAI capabilities can enhance electronic health record (EHR) systems, streamline workflows, support documentation quality initiatives, and ameliorate patient engagement processes. Caution is warranted, however, as these systems are not without risk. Case studies that highlight early experiences are detailed in the sections that follow.

Generative AI is being adopted to enhance clinical documentation in multiple ways. One use case addresses the time-consuming task of coding surgical procedures. Clinicians typically dictate an operative report that undergoes coding and reconciliation by trained medical coders. These coders use human judgment to ensure that the proper procedure and diagnosis codes are applied, but coders sometimes miss important clinical details because the reports lack careful organization and formatting. An AI-based solution assists the coders by automatically summarizing the report, including collapsing multi-part sections, providing bulleted lists of critical details, and normalizing terminology for improved readability. Preliminary evaluation suggests that AI augmentation may enhance coder productivity.



Equation 3: Training objective: Cross-Entropy / Negative Log-Likelihood



Let the true next token at step t be y_t (one-hot). The probability assigned to the correct token is:

$$p_t = P(x_t = y_t \mid x_{<t})$$

Step 1: likelihood of the correct token

$$\mathcal{L}_t = p_t$$

Step 2: log-likelihood

$$\log \mathcal{L}_t = \log(p_t)$$

Step 3: negative log-likelihood (loss)

$$\ell_t = -\log(p_t)$$

Step 4: average over all tokens

$$\mathcal{L} = \frac{1}{T} \sum_{t=1}^T -\log(P(x_t \mid x_{<t}))$$

6.1. Electronic Health Record Enhancements

Deployment of Generative AI capabilities in clinical settings has generated an array of productive and efficient use cases that provide critical, augmented functionality for electronic health record systems. The breadth of effort is illustrated by three initiatives, which span the use of AI-native transcription for converting audio of clinical encounters directly into the EHR; improved note-taking with clinically-situated support for Completion, Summarization, and Redaction; and a dedicated exploration of methods for quantifying content quality.

First, the integration of large language models with large-scale automatic speech recognition and even more massive audio captioning models can transform spoken audio from clinical conversations into textual transcripts that are directly entered into the EHR. While AI-native transcription supports clinicians in documenting encounters, it also improves patient-centredness by restoring the natural dynamic of dialogue—the back-and-forth pattern of speech, reflection, and clarification that characterizes human conversation. Conversations are dynamic, spontaneous, and directional, shifting ownership

of speaking from one party to the other; unlike standard dictation, patients are not constrained by the accuracy of the recognition because they are not reading; meaning is conveyed by the delivery, and it is clear when they are done with a point—it is not possible to momentarily scuttle a rambling response. Heterogeneity and disfluent speech are to be expected, but in the end it is the effectiveness of the interaction that matters.

7. Conclusion

Generative AI should be described in formal, concise terms with precise terminology, focusing on healthcare applications, ethical considerations, and measurable outcomes, using a professional tone consistent with Clinical Documentation and Patient Engagement 2024. Key findings must be summarized, limitations acknowledged, future directions proposed, and ethical, legal, and patient-centered considerations emphasized.

Generative AI describes a category of artificial intelligence technologies capable of automatically producing media or information. In the healthcare domain, applications are emerging for clinical documentation and patient interaction. Automated clinical notes enable a significant reduction in human labor by tailoring note generation to each provider's clinical, cognitive, and information-seeking needs. Patient-facing solutions contribute to both education and administrative needs, addressing health literacy and cultural competence considerations. Use cases include scheduling and triage. Nevertheless, risk mitigation remains critically important; generative AI systems may misrepresent facts, invent responses, present information at inappropriate levels of complexity, or infer sensitive information not disclosed by patients themselves.

To deploy these solutions successfully, healthcare organizations should assess vendor capabilities in key areas: interoperability, data provenance and governance, data privacy and security, harmonization of third-party risk management processes, change management, and clinician adoption. Although these solutions require further validation and development, careful evaluation and management make them immediately feasible, making a



timely impact on clinical documentation and patient engagement workloads.




Fig 4: Clinical Deployment Framework Maturity

7.1. Final Thoughts and Future Directions in Healthcare AI

As generative AI's potential becomes increasingly clear, the most pressing question is not its power but how it will improve care delivery and patient outcomes. Healthcare organizations should focus on use cases that can be operationalized quickly, with safety and efficacy documented and demonstrated. Vendors can help by prioritizing capabilities that integrate with existing infrastructure rather than replacing solutions that are already working.

Generative AI is still in its infancy, and healthcare organizations should collaborate with manufacturers and the broader ecosystem to develop industry standards. These standards will ensure safe, secure, and ethical solutions. It is essential that organizations assess products thoroughly to mitigate risk and avoid vendor lock-in. Data security and patient consent are paramount, and an updated data governance framework should prioritise these qualities. Patient engagement with AI, as with all other technologies, needs to be led from within. Decision-making must be clear and transparent, user experience simple and straightforward. Redesign using AI should occur hand-in-hand with comprehensive training and usability testing to help clinicians feel empowered rather than threatened. Training should also be combined with clearly defined change-management initiatives, support and encouragement from leadership, and a focus on clinical

relevance. Performance metrics should be defined ahead of deployment to provide clinicians with feedback on how AI is transforming their work. Organizations must test, evaluate, learn and adapt.

Generative AI solutions for clinical documentation and patient engagement have made great strides and are already being tested and evaluated in real-world deployments. Initial indications are positive, but implementation must be carefully managed, development decisions must be transparent, products must be rigorously evaluated, and the fundamentals of patient-centred care must remain uppermost in mind.

8. References

- [1] Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
- [2] Shortliffe, E. H., & Sepúlveda, M. J. (2018). Clinical decision support in the era of artificial intelligence. *JAMA*, 320(21), 2199–2200.
- [3] Sutton, R. T., Pincock, D., Baumgart, D. C., Sadowski, D. C., Fedorak, R. N., & Kroeker, K. I. (2020). An overview of clinical decision support systems. *Journal of Medical Internet Research*, 22(10), e24862.
- [4] Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., & Wang, Y. (2017). Artificial intelligence in healthcare. *Stroke and Vascular Neurology*, 2(4), 230–243.
- [5] Bates, D. W., Saria, S., Ohno-Machado, L., Shah, A., & Escobar, G. (2014). Big data in health care. *Health Affairs*, 33(7), 1123–1131.



- [6] Musen, M. A., Middleton, B., & Greenes, R. A. (2014). Clinical decision-support systems. In *Biomedical informatics* (pp. 643–674). Springer.
- [7] Greenes, R. A. (2014). *Clinical decision support: The road ahead*. Academic Press.
- [8] Sendak, M. P., D'Arcy, J., Kashyap, S., Gao, M., Nichols, M., Corey, K., Ratliff, W., & Balu, S. (2020). A path for translation of machine learning products into healthcare delivery. *EMJ Innovations*, 4(1), 70–80.
- [9] Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P., Thadaney-Israni, S., & Goldenberg, A. (2019). Do no harm. *Nature Medicine*, 25(9), 1337–1340.
- [10] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Proceedings of the ACM SIGKDD Conference, 1135–1144.
- [11] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv Preprint*.
- [12] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Pedreschi, D., & Giannotti, F. (2019). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1–42.
- [13] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., et al. (2020). Explainable artificial intelligence. *Information Fusion*, 58, 82–115.
- [14] Shneiderman, B. (2020). Human-centered artificial intelligence. *International Journal of Human-Computer Interaction*, 36(6), 495–504.
- [15] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms. *Big Data & Society*, 3(2), 1–21.
- [16] Floridi, L., Cows, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707.
- [17] Holmes, J. H., Durbin, D. R., & Glass, M. (2021). Harnessing artificial intelligence to improve patient care. *JAMA*, 325(12), 1177–1178.
- [18] Elwyn, G., Frosch, D., Thomson, R., Joseph-Williams, N., Lloyd, A., Kinnersley, P., Cording, E., Tomson, D., Dodd, C., Rollnick, S., Edwards, A., & Barry, M. (2012). Shared decision making. *Journal of General Internal Medicine*, 27(10), 1361–1367.
- [19] Ammenwerth, E., Iller, C., & Mahler, C. (2008). IT adoption and the interaction of task, technology and individuals. *Methods of Information in Medicine*, 47(3), 245–252.
- [20] DeLone, W. H., & McLean, E. R. (2003). The DeLone and McLean IS success model. *Journal of Management Information Systems*, 19(4), 9–30.
- [21] Parasuraman, A., Zeithaml, V. A., & Berry, L. L. (1988). SERVQUAL. *Journal of Retailing*, 64(1), 12–40.
- [22] Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work.



Academy of Management Annals, 14(1), 366–410.

[23] Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management. *Academy of Management Review*, 46(1), 192–217.

[24] van der Aalst, W. (2021). Process mining and real-time analytics. *Communications of the ACM*, 64(8), 76–83.

[25] Jans, M., Alles, M., & Vasarhelyi, M. (2014). Process mining in auditing and healthcare. *Accounting Horizons*, 28(4), 815–843.

[26] Vial, G. (2019). Understanding digital transformation. *MIS Quarterly*, 43(1), 223–247.

[27] Holmström, J., Holweg, M., Lawson, B., Pil, F., & Wagner, S. (2019). The digitalization of operations. *Journal of Operations Management*, 65(8), 728–734.

[28] Kearns, M., & Roth, A. (2020). *The ethical algorithm*. Oxford University Press.

[29] Molnar, C. (2022). *Interpretable machine learning* (2nd ed.). Leanpub.

[30] OECD. (2022). *Trustworthy artificial intelligence in the health sector*. OECD Publishing.

[31] World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. WHO Press.

[32] NIST. (2023). *Artificial intelligence risk management framework*. National Institute of Standards and Technology.

[33] Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209.

[34] Liu, X., Glocker, B., McCradden, M. M., Ghassemi, M., Denniston, A. K., & Oakden-Rayner, L. (2022). The medical algorithmic audit. *Nature Medicine*, 28(7), 1333–1343.

[35] Lee, P., Bubeck, S., & Petro, J. (2023). Benefits, limits, and risks of GPT-4 as an AI assistant for medicine. *New England Journal of Medicine*, 388(13), 1233–1239.