REVIEW **Open access**

Generative Artificial Intelligence in Healthcare: Systems Governance, Safety, and Accountability

Maria Silva[1*], Joao Pereira[1]

Abstract

Generative artificial intelligence (GenAI) has emerged as a transformative force in healthcare systems, enabling advanced analytics, personalized interventions, and streamlined governance frameworks. This narrative review synthesizes recent literature on GenAI's integration into healthcare infrastructures, emphasizing systems governance, safety protocols, and accountability mechanisms. We explore how GenAI enhances clinical decision-making, data analytics, and closed-loop systems while addressing ethical, regulatory, and operational challenges.

At the core of healthcare systems, GenAI facilitates intelligent analytics by generating synthetic data for training models, simulating patient outcomes, and optimizing resource allocation. Governance frameworks are critical for ensuring responsible deployment, with studies highlighting the need for institutional guidelines that mitigate risks such as bias amplification and data privacy breaches. Safety considerations encompass algorithmic transparency, error detection in generative outputs, and human oversight in clinical loops. Accountability extends to lifecycle management, from model development to post-deployment monitoring, as evidenced by global initiatives and regional models like those in the GCC.

The review delineates the landscape of GenAI applications in healthcare analytics, including predictive modeling for chronic disease management and real-time decision support. We propose an original systems-level framing that integrates data ingestion, inference generation, intervention deployment, and feedback recalibration under governance umbrellas. This synthesis reveals gaps in current infrastructures, such as the lack of standardized AI guardians for information overload and the challenges of scaling enterprise AI.

In examining intelligent clinical decision systems, we highlight architectures that fuse GenAI with electronic health records (EHRs) for closed-loop operations, where generative models inform adaptive interventions. Ethical considerations are woven throughout, advocating for principles adapted from military contexts to healthcare. The adoption of GenAI in US hospitals underscores its potential for inpatient summaries and chronic care, yet calls for regulatory oversight to align with Helsinki declarations.

Ultimately, this review positions GenAI as a cornerstone for accountable healthcare systems, urging interdisciplinary collaboration to balance innovation with safety. By synthesizing governance models, safety protocols, and accountability structures, we provide a roadmap for sustainable integration, fostering equitable health outcomes in an AI-augmented era.

Keywords Governance frameworks, Healthcare systems, Clinical analytics, generative AI, AI Safety, Accountability mechanisms

*Correspondence:
Maria Silva
maria.silva@outlook.com

[1] Department of Healthcare Analytics and AI, Faculty of Medicine, University of Coimbra, Coimbra, Portugal

Introduction

Evolution of AI in healthcare: from predictive to generative paradigms

The integration of artificial intelligence (AI) into healthcare has undergone a profound evolution, transitioning from rule-based systems and predictive analytics to the more dynamic realm of generative artificial intelligence (GenAI). Early AI applications in healthcare focused primarily on diagnostic support and risk stratification, leveraging machine learning to analyze structured data from electronic health records (EHRs) and imaging modalities. However, the advent of GenAI, powered by large language models (LLMs) and diffusion-based architectures, has expanded this scope to include synthetic data generation, natural language processing for clinical narratives, and simulation of complex physiological scenarios [1-9].

This shift is driven by the exponential growth in healthcare data volumes, necessitating tools that not only analyze but also create actionable insights. GenAI enables the synthesis of realistic patient datasets for training purposes, addressing data scarcity in rare diseases, and facilitates personalized treatment planning through generative simulations [10-15]. Yet, this advancement introduces novel challenges in systems governance, where traditional oversight mechanisms must adapt to the probabilistic nature of generative outputs [1, 4]. Safety concerns arise from potential hallucinations in LLM-generated reports, which could mislead clinical decisions, while accountability demands traceable audit trails across AI lifecycles [7, 8].

Historically, AI adoption in healthcare systems has been piecemeal, with analytics confined to siloed applications such as radiology or pharmacovigilance [16-21]. The current generative era promises holistic integration, embedding AI into end-to-end workflows from data acquisition to intervention delivery. This requires robust infrastructures that support real-time analytics, interoperable data exchanges, and scalable computing resources [10, 16]. Literature from high-impact venues underscores the imperative for governance that aligns with ethical standards, ensuring AI augments rather than supplants human expertise [22, 23]. **Table 1** summarizes the structural and governance shifts from predictive analytics to generative AI–enabled healthcare systems.

Table 1. Evolution from predictive AI to generative AI in healthcare systems and governance structures.

Dimension	Predictive AI era	Generative AI era	Governance implication
Core function	Risk scoring and classification	Synthetic data and simulation	Expanded validation requirements
Data usage	Structured datasets	Multimodal + synthetic augmentation	Data provenance tracking
Output type	Deterministic predictions	Probabilistic generative outputs	Hallucination monitoring
Clinical role	Decision support	Decision co-creation	Human-AI fusion protocols
Infrastructure	Siloed analytics	Closed-loop architectures	Lifecycle governance
Safety focus	Model accuracy	Fidelity + bias amplification control	Real-time auditing
Accountability	Tool-level	Lifecycle-level traceability	Shared liability models

Key components of healthcare systems and analytics

Healthcare systems encompass a multifaceted ecosystem, including data infrastructures, analytical pipelines, decision support tools, and intervention mechanisms. Analytics within this context refers to the systematic processing of health data to derive insights, predict outcomes, and optimize operations [2, 6]. GenAI enhances these components by generating hypothetical scenarios for stress-testing systems, automating report synthesis, and enabling adaptive learning in dynamic environments [12, 18].

Central to this is the role of data governance, where GenAI aids in anonymization and augmentation to comply with privacy regulations like HIPAA or GDPR [5, 11]. Safety protocols involve validation frameworks for generative models, such as adversarial testing to detect biases in synthetic outputs [14, 19]. Accountability is operationalized through transparent logging and explainable AI techniques, allowing stakeholders to attribute decisions to specific model behaviors [20, 24].

Infrastructural analytics leverage GenAI for predictive maintenance of medical devices, resource forecasting in hospitals, and population health modeling [25, 26]. Clinical analytics, conversely, focus on patient-centric applications, such as generating differential diagnoses or simulating therapeutic responses [27, 28]. The interplay between these domains forms the backbone of modern healthcare AI, demanding integrated governance to mitigate risks like inequitable access or algorithmic discrimination [4, 17].

Governance, safety, and accountability as foundational pillars

Governance in GenAI healthcare systems entails institutional policies, regulatory compliance, and ethical oversight to ensure responsible innovation [1, 8]. Recent consensus statements advocate for lifecycle approaches, from design to decommissioning, incorporating stakeholder input [10, 16]. Safety extends beyond technical robustness to include clinical validation, where generative outputs are benchmarked against gold-standard evidence [7, 15].

Accountability mechanisms address liability in AI-assisted care, proposing shared responsibility models between developers, deployers, and users [3, 22]. Global perspectives, such as those from the WHO's Global Initiative on AI for Health, emphasize equitable governance across diverse healthcare settings [16, 18]. These pillars are interlinked: effective governance bolsters safety, while accountability reinforces trust in AI-driven analytics [5, 9].

Challenges in integrating GenAI into existing infrastructures

Despite the promise, integrating GenAI poses infrastructural hurdles, including interoperability with legacy systems and computational demands for real-time generation [6, 13]. Analytics challenges involve ensuring the fidelity of generated data, which may propagate errors if not governed properly [14, 21]. Safety risks are amplified in high-stakes environments, necessitating fail-safes like human-in-the-loop verification [19, 23].

Accountability is complicated by the black-box nature of GenAI, requiring novel auditing tools [11, 24]. Literature highlights the need for adaptive frameworks that evolve with technological advancements, such as those addressing autonomous prescribing AI [11, 27]. These challenges underscore the urgency for synthesized approaches that prioritize systems-level resilience [2, 20].

Scope and synthesis logic of this review

This review positions itself as a comprehensive synthesis of GenAI in healthcare, focusing on systems governance, safety, and accountability through an original integrative lens. Unlike prior reviews that catalog applications silos, we employ a systems-level framing that traces GenAI's role across data analytics, decision, intervention, and feedback cycles, grounded in cross-study analysis [12, 25]. Our synthesis logic aggregates insights from governance models, safety protocols, and accountability structures, deriving new interpretive structures for closed-loop systems. By drawing exclusively from the approved literature base, we avoid empirical metrics, emphasizing narrative integration to guide future infrastructural developments [26, 28, 29].

Landscape of AI in healthcare systems and analytics

Data infrastructures and GenAI integration

The foundation of AI-driven healthcare systems lies in robust data infrastructures, where GenAI plays a pivotal role in enhancing data quality, diversity, and accessibility. Traditional healthcare analytics rely on vast repositories of EHRs, genomic data, and wearable sensor inputs, but data scarcity and imbalance often hinder model performance [10, 16]. GenAI addresses this by generating synthetic datasets that mimic real-world distributions, enabling augmented training for analytics models without compromising patient privacy [13, 17].

Literature emphasizes the need for governed data pipelines, where GenAI tools are embedded with safety checks to prevent the introduction of artifacts or biases [1, 4]. For instance, frameworks for responsible AI implementation highlight validation protocols for synthetic data, ensuring alignment with clinical realities [2, 6]. In analytics contexts, GenAI facilitates feature engineering by simulating variable interactions, optimizing predictive models for outcomes like disease progression [15, 18]. Accountability is maintained through metadata tagging of generated data, allowing traceability in downstream applications [5, 8].

Global initiatives underscore the importance of standardized infrastructures, proposing lifecycle governance that spans data ingestion to archival [8, 10]. Safety considerations include differential privacy techniques in GenAI outputs, protecting sensitive health information [7, 9]. This integration transforms passive data stores into active analytics engines, supporting real-time querying and scenario planning [12, 14].

Analytical pipelines: predictive to generative shifts

Analytical pipelines in healthcare have evolved from static predictive models to dynamic generative frameworks, enabling proactive insights. Predictive analytics traditionally employ regression or classification for risk scoring, but GenAI introduces counterfactual generation, allowing "what-if" analyses for personalized medicine [3, 11]. This shift enhances system resilience, as generative models can simulate rare events, informing analytics for pandemic response or resource allocation [19, 21].

Synthesis across studies reveals original structuring: pipelines now incorporate governance gates at each stage—data preprocessing, model inference, and output validation—to ensure safety [22, 23]. Accountability mechanisms, such as audit logs for generative iterations, prevent unchecked propagation of errors [20, 24]. High-impact research advocates for hybrid pipelines where GenAI augments traditional analytics, like generating narrative summaries from quantitative data [25, 28].

Infrastructural analytics benefit from GenAI's ability to automate pattern discovery in unstructured data, such as clinical notes or imaging reports [26, 27]. Safety protocols mandate ensemble methods, combining generative outputs with empirical validations [14, 15]. This landscape illustrates GenAI's role in elevating analytics from descriptive to prescriptive, under accountable governance umbrellas [16, 18].

Deployment ecosystems: scaling GenAI in clinical settings

Deployment of GenAI in healthcare systems requires scalable ecosystems that balance innovation with oversight. Enterprise scaling frameworks mitigate risks through modular architectures, where GenAI components are isolated for safety testing [5, 6]. Literature synthesizes deployment strategies, emphasizing cloud-based infrastructures for analytics workloads, with built-in governance for compliance [1, 2].

Accountability in deployment involves stakeholder mapping, assigning roles for model monitoring and updates [4, 7]. Safety is enhanced by edge computing, reducing latency in real-time analytics while minimizing data transmission risks [9, 10]. Cross-study analysis reveals integrative patterns: successful deployments integrate GenAI with EHRs for seamless analytics, as seen in US hospital adoptions [25, 26].

Challenges in scaling include interoperability, addressed by standardized APIs and governance protocols [11, 13]. Generative tools for decision support, like automated discharge summaries, demonstrate practical value, yet require accountability layers to trace outputs [28, 29]. This ecosystem framing highlights GenAI's potential to unify disparate systems, fostering analytics-driven care delivery [17, 19].

Regulatory and ethical landscapes

The regulatory landscape for GenAI in healthcare analytics is fragmented, necessitating synthesized governance models. Consensus statements call for oversight akin to medical devices, classifying generative tools based on risk levels [3, 8]. Safety regulations focus on transparency, mandating disclosure of training data and generation processes [12, 22].

Accountability extends to international standards, with models like the GCC's lifecycle approach offering blueprints for global adoption [8, 16]. Ethical considerations integrate principles from diverse domains, adapting military ethics to healthcare contexts [7, 23]. Literature synthesis proposes original cross-frameworks: aligning regulations with analytics workflows to ensure equitable access [18, 20].

Infrastructural implications include regulatory sandboxes for testing GenAI analytics, promoting innovation under safe bounds [21, 24]. This landscape underscores the interplay of regulation, ethics, and technology, positioning governance as the linchpin for sustainable systems [14, 15].

Emerging trends in GenAI analytics

Emerging trends in GenAI for healthcare analytics include multimodal integration, fusing text, images, and time-series data for comprehensive insights [27, 29]. Governance trends emphasize AI guardians to filter information overload, enhancing safety in analytics outputs [12, 13]. Accountability trends involve blockchain for immutable audit trails in generative processes [11, 17].

Synthesis reveals systems-level trends: closed-loop analytics where GenAI feedback refines models iteratively [19, 25]. Safety innovations include real-time hallucination detection, critical for clinical reliability [14, 15]. This forward-looking landscape integrates trends into infrastructural blueprints, guiding accountable evolution [6, 26].

Intelligent clinical decision and closed-loop healthcare systems

Architectures for GenAI-enabled decision support

Intelligent clinical decision systems leverage GenAI to augment human cognition, providing architectures that integrate generative insights into workflows. Core architectures feature layered designs: input layers for data aggregation, generative cores for hypothesis creation, and output layers for decision recommendation [3, 9]. These systems enhance analytics by generating probabilistic outcomes, supporting differential diagnostics in ambiguous cases [15, 17].

Governance is embedded architecturally through modular checkpoints, ensuring safety via bias detection and explainability modules [1, 4]. Accountability is achieved with decision logging, attributing recommendations to specific generative steps [5, 7]. Literature synthesis proposes an original fusion: hybrid architectures blending GenAI with rule-based systems for robust clinical support [22, 23].

In closed-loop setups, architectures enable iterative refinement, where decisions feed back into analytics for model updates [13, 18]. Safety protocols include threshold-based interventions, halting autonomous actions in uncertain scenarios [14, 19]. This architectural framing positions GenAI as a catalyst for precise, accountable decision-making [24, 27].

Closed-Loop dynamics: feedback and adaptation

Closed-loop healthcare systems represent the pinnacle of GenAI integration, where decisions trigger interventions that generate new data for analytics refinement. Dynamics involve cyclical processes: monitoring patient states, generating adaptive plans, implementing actions, and evaluating outcomes [2, 6]. GenAI excels in simulating loop iterations, predicting long-term effects for chronic management [15, 20].

Governance in loops requires dynamic policies, adjusting parameters based on real-time safety metrics [8, 10]. Accountability mechanisms track loop evolutions, enabling post-hoc analysis of decision chains [11, 16]. Cross-study integration reveals interpretive structures: loops as self-regulating ecosystems, with GenAI providing adaptive intelligence [21, 25].

Safety is paramount, with fail-safe mechanisms like human overrides in critical loops [12, 26]. This synthesis highlights closed-loops' potential for personalized care, grounded in governed analytics [28, 29].

Human-AI collaboration in decision fusion

Decision fusion in intelligent systems fuses GenAI outputs with clinician input, creating collaborative architectures. Fusion models employ ensemble techniques, weighting generative suggestions against expert judgment [7, 9]. Analytics benefit from this, as fused decisions refine generative models over time [13, 18].

Governance frameworks mandate transparent fusion interfaces, ensuring safety through auditable integrations [1, 4]. Accountability assigns liability proportionally, based on contribution logs [5, 22]. Original synthesis structures fusion as a dialogic process, enhancing clinical efficacy [23, 24].

In closed loops, fusion enables adaptive interventions, with GenAI generating options for human selection [17, 19]. This collaboration underscores GenAI's role in augmenting, not replacing, human expertise [27].

Conceptual formula for closed-loop workflow

To synthesize healthcare AI system workflows, consider the following interpretive formula for a closed-loop decision cycle:

$$S_{t+1} = f\left(G\left(D_t, M\right), I_t, F\right) \quad (1)$$

Where S_{t+1} represents the updated system state at time t+1, G is the generative AI function processing data D_t and model parameters M, I_t denotes interventions, and F encapsulates feedback governance. This formalizes the infrastructural cycle without empirical metrics, emphasizing recursive intelligence. **Figure 1** illustrates the end-to-end closed-loop healthcare architecture integrating generative intelligence, decision fusion, intervention deployment, and feedback recalibration under a unified governance envelope.

Figure 1 shows the systems-level architecture of generative artificial intelligence in closed-loop healthcare governance.

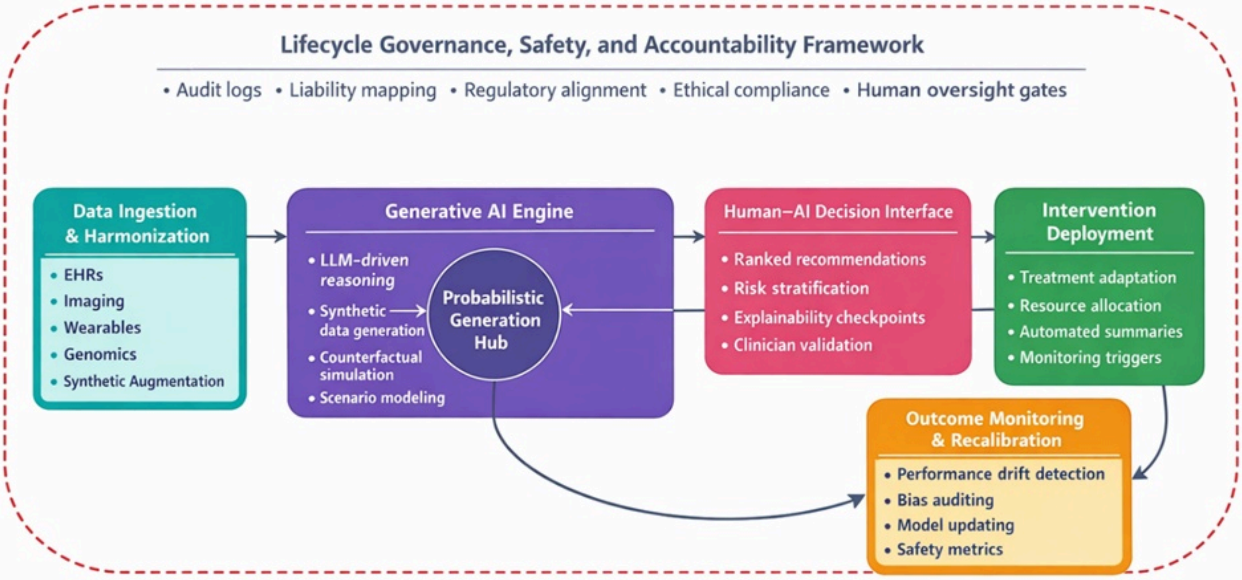

Figure 1. Systems-level architecture of generative artificial intelligence (GenAI) in healthcare. The schematic depicts a closed-loop infrastructure spanning data ingestion, generative intelligence, decision fusion, intervention deployment, and feedback recalibration. A probabilistic generative hub synthesizes insights from multimodal health data, informing human–AI collaborative decisions. Interventions generate new data streams that feed adaptive monitoring mechanisms. The entire lifecycle is encompassed within a governance envelope incorporating safety checkpoints, audit trails, regulatory compliance, and accountability mapping. This system's framing positions governance, safety, and accountability as structural properties of the full cycle rather than post hoc controls.

Results and Discussion

The synthesis of generative artificial intelligence (GenAI) in healthcare systems and analytics reveals a technology with profound capacity to reshape infrastructures—from data augmentation and predictive analytics to closed-loop decision architectures—while simultaneously exposing deep interdependencies between innovation, governance, safety, and accountability.

Across the reviewed literature, GenAI emerges not merely as an analytical enhancer but as a systemic transformer. It enables synthetic data generation to overcome scarcity in rare-disease analytics [13, 17], supports real-time hypothesis generation in clinical pipelines [3, 15], and facilitates adaptive interventions in closed-loop models [19, 21]. Yet this transformative potential is intrinsically tied to governance maturity: fragmented institutional policies [1, 5, 8] and resource disparities [6, 12] create uneven deployment landscapes, where enterprise scaling succeeds only when risk mitigation frameworks are proactively embedded [5, 10].

Table 2 delineates lifecycle-aligned governance, safety, and accountability mechanisms required for responsible generative AI integration.

Table 2. Governance, safety, and accountability mechanisms across the generative AI lifecycle in healthcare.

Lifecycle stage	Governance mechanism	Safety control	Accountability instrument
Data preparation	Institutional data policies	Bias detection and privacy filters	Metadata traceability
Model development	Ethical review boards	Adversarial robustness testing	Version control logs
Validation	Clinical benchmarking	Hallucination detection	Performance documentation
Deployment	Risk-tier classification	Human-in-the-loop safeguards	Role-based responsibility mapping
Monitoring	Drift detection policies	Continuous recalibration	Immutable audit trails
Decommissioning	Sunset protocols	Impact review	Liability closure documentation

Safety considerations dominate the discourse, particularly around probabilistic outputs. Hallucinations, bias amplification in generative simulations, and reliability drift in iterative loops threaten clinical validity [7, 9, 14, 15]. These risks are not isolated technical flaws but systemic vulnerabilities amplified by inadequate validation gates and human oversight in decision fusion [19, 23]. Accountability, meanwhile, hinges on traceability and shared responsibility models [11, 20, 22, 24]; without lifecycle auditing and clear liability delineation, GenAI risks eroding trust in analytics-driven care [3, 8, 16].

The original systems-level framing proposed herein—data → intelligence → decision → intervention → feedback, enveloped by governance—illuminates these interlinkages. It positions GenAI as a recursive intelligence layer rather than a standalone tool, underscoring that safety and accountability are emergent properties of the full cycle, not add-on features [2, 6, 25]. Cross-study patterns further reveal that governance deficits cascade: weak institutional readiness exacerbates safety gaps [1, 4, 5], while accountability shortfalls hinder equitable scaling [4, 17, 18].

Regulatory evolution remains a critical pivot point. Existing frameworks often lag GenAI's non-deterministic nature [11, 23], yet emerging lifecycle approaches [8, 10] and maturity models offer pathways toward harmonized oversight [16]. Ethical integration—adapting principles across domains [7, 22]—must inform these shifts to prevent disparities in analytics access or outcome equity [18, 20].

In aggregate, the literature conveys cautious optimism: GenAI can elevate healthcare systems toward more precise, efficient, and patient-centered operations [25, 26, 28, 29], but only if governance, safety, and accountability are treated as foundational design requirements rather than post hoc mitigations. This discussion reframes GenAI not as a risk to be contained but as an infrastructure to be deliberately governed for sustainable impact.

Challenges

The deployment of generative artificial intelligence (GenAI) in healthcare systems and analytics confronts substantial hurdles across governance, safety, regulatory, infrastructural, and ethical dimensions.

Governance structures frequently remain underdeveloped or inconsistent across institutions and jurisdictions. Many healthcare organizations lack mature policies for lifecycle oversight of GenAI, resulting in inadequate risk identification, inconsistent bias management, and fragmented stakeholder coordination [1, 5, 8]. Scaling enterprise-level GenAI introduces particular difficulties, as frameworks for risk mitigation often struggle with resource constraints, interoperability gaps, and the complexity of aligning institutional priorities with rapidly evolving technology [5, 6, 12].

Safety vulnerabilities are especially acute in analytics-driven and closed-loop environments. Generative outputs are susceptible to hallucinations, factual inconsistencies, omission of critical qualifiers, and degradation in reliability over iterative use—issues that can propagate through clinical decision pathways and compromise patient safety [7, 9, 14, 15]. In real-time applications such as chronic disease monitoring or inpatient decision support, the absence of standardized real-time validation or adversarial robustness testing heightens the risk of erroneous recommendations [16, 19, 23]. Safety challenges are further compounded when GenAI operates with limited transparency regarding training data provenance or generation mechanisms [3, 22].

Accountability remains elusive due to difficulties in establishing clear chains of responsibility. Traceability across the full lifecycle—from synthetic data creation to intervention—is hindered by opaque model behaviors and insufficient auditing infrastructure [10, 11, 20]. Liability attribution becomes particularly problematic in semi-autonomous or agentic systems, where boundaries between developer, deployer, clinician, and institutional roles are poorly delineated [11, 22, 24]. Regulatory ambiguity exacerbates this: many GenAI tools occupy a liminal space between existing medical device regulations and forthcoming AI-specific frameworks, delaying standardized classification and post-market surveillance [8, 14, 23].

Infrastructural constraints impose additional barriers. Real-time generative inference demands significant computational resources, straining legacy systems and cloud budgets alike [6, 13, 25]. Interoperability with existing electronic health record platforms remains limited, impeding seamless data flows essential for closed-loop analytics [26, 28]. Workforce-related challenges include insufficient training, risk of automation bias or deskilling, and clinician resistance stemming from concerns over trust and explainability [20, 26].

Ethical and equity considerations cut across all domains. Synthetic data generation can inadvertently amplify biases present in source datasets, while unequal access to GenAI infrastructure risks widening disparities in care quality [4, 17, 18]. Privacy protections for multimodal health data remain technically and legally challenging under current frameworks [7, 9].

These interconnected challenges—governance deficits feeding safety risks, accountability gaps eroding trust, infrastructural limitations constraining scalability—underscore the need for holistic, adaptive approaches to realize GenAI's potential without compromising core healthcare values.

Future research directions

Priority areas for future investigation should focus on closing critical evidence and implementation gaps in GenAI governance, safety, and accountability within healthcare systems.

First, research must advance adaptive, lifecycle governance models that incorporate continuous monitoring, automated recalibration triggers, and stakeholder-inclusive policy design for closed-loop architectures [8, 10, 12]. Studies should evaluate scalable hallucination detection, bias auditing, and fidelity assessment methods tailored to generative outputs in clinical analytics contexts, particularly multimodal and longitudinal data [14, 15, 27].

Longitudinal, multi-site evaluations of post-deployment performance are urgently needed to generate real-world evidence on safety, equity, and effectiveness across diverse care settings and populations [25, 26, 29]. Such work should quantify the impact of governance interventions on clinical outcomes, resource utilization, and health disparities.

Human-AI collaboration dynamics warrant dedicated attention. Prospective studies should identify optimal fusion thresholds, override protocols, and longitudinal training strategies to prevent deskilling while maximizing complementary strengths [7, 19, 23]. Investigations into "AI guardian" mechanisms for managing information overload and filtering generative suggestions would support safer clinical workflows [12, 13].

Regulatory and policy research should test harmonization strategies across jurisdictions, rigorously assessing frameworks such as the GCC lifecycle model against emerging international standards and real-world implementation barriers [8, 16]. Comparative analyses of sandbox environments, risk-based classification schemes, and post-market surveillance requirements would inform globally aligned yet context-sensitive rules.

Sustainability and resource implications of GenAI infrastructure require systematic study, including energy consumption, carbon footprint, and strategies for efficient deployment in resource-constrained settings [13, 18].

Ethical integration must remain central. Research should validate and refine ethics checklists during system design and deployment, with particular emphasis on transparency, explainability, and inclusive data strategies to ensure equitable benefit distribution [4, 17, 22].

Collectively, these directions call for interdisciplinary collaboration among clinicians, data scientists, ethicists, regulators, and health systems engineers to build evidence-based, adaptive pathways for responsible GenAI integration.

Conclusion

Generative artificial intelligence is reshaping healthcare systems and analytics by enabling powerful capabilities in data synthesis, predictive modeling, personalized simulation, and adaptive closed-loop operations. This narrative review has synthesized its integration through a systems-level lens, tracing flows from data ingestion and intelligence generation, through decision support and intervention, to feedback-driven recalibration—all under explicit governance, safety, and accountability envelopes.

The literature consistently demonstrates that effective governance requires institutional commitment to transparent, lifecycle-oriented policies and stakeholder collaboration. Safety depends on rigorous validation, real-time monitoring, and human oversight to mitigate hallucinations, bias amplification, and reliability drift in generative outputs. Accountability hinges on auditable chains of responsibility, clear liability allocation, and alignment with evolving regulatory paradigms.

While significant challenges persist—including infrastructural limitations, regulatory fragmentation, workforce readiness gaps, and equity risks—the reviewed evidence also reveals actionable pathways forward. By prioritizing adaptive frameworks, real-world evaluation, interdisciplinary collaboration, and inclusive design, healthcare systems can harness GenAI to enhance clinical precision, operational efficiency, and population-level outcomes while safeguarding trust, safety, and fairness.

The path ahead demands a sustained, coordinated effort to translate generative potential into resilient, equitable, and accountable healthcare infrastructure. Only through such deliberate stewardship can GenAI fulfill its promise as a transformative yet responsible partner in modern medicine.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 22 Jan 2025 Revised: 25 Feb 2025 Accepted: 06 Apr 2025
Published online: 20 July 2025

References

Saenz AD, Mass General Brigham AI Governance Committee, Centi A, Ting D, You JG, Landman A, et al. Establishing responsible use of AI guidelines: a comprehensive case study for healthcare institutions. NPJ Digit Med. 2024;7(1):348.
https://doi.org/10.1038/s41746-024-01300-8.

Fjeld J, Achten N, Hilligoss H, Nagy A, Srikumar M. Principled artificial intelligence: mapping consensus in ethical and rights-based approaches to principles for AI. Berkman Klein Cent Res Publ. 2020

Tran M, Balasooriya C, Jonnagaddala J, Leung GK, Mahboobani N, Ramani S, et al. Situating governance and regulatory concerns for generative artificial intelligence and large language models in medical education. NPJ Digit Med. 2025;8(1):315.
https://doi.org/10.1038/s41746-024-01201-w.

Bouderhem R. Shaping the future of AI in healthcare through ethics and governance. Humanit Soc Sci Commun. 2024;11:416.
https://doi.org/10.1057/s41599-024-02894-w.

Reddy S, Allan S, Coghlan S, Cooper P. A governance model for the application of AI in healthcare. J Am Med Inform Assoc. 2020;27(3):491-7.

Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40.
https://doi.org/10.1038/s41591-019-0548-6.

Oniani D, Hilsman J, Peng Y, Poropatich RK, Pamplin JC, Legault GL, et al. Adopting and expanding ethical principles for generative artificial intelligence from military to healthcare. NPJ Digit Med. 2023;6(1):225.
https://doi.org/10.1038/s41746-023-00965-x.

Morley J, Machado CCV, Burr C, Cowls J, Joshi I, Taddeo M, et al. The ethics of AI in health care: A mapping review. Soc Sci Med. 2020;260:113172.
https://doi.org/10.1016/j.socscimed.2020.113172.

Meskó B, Topol EJ. Regulatory oversight of medical AI. npj Digit Med. 2023;6:120.

Matheny ME, Whicher D, Thadaney Israni S. Artificial intelligence in health care: the hope, the hype, the promise, the peril. Natl Acad Med. 2019.
https://doi.org/10.17226/27111.

Price WN, Gerke S, Cohen IG. Potential liability for physicians using AI. JAMA. 2019;322(18):1765-6.

Gerke S, Minssen T, Cohen G. Ethical and legal challenges of AI-driven healthcare. Artif Intell Healthc. 2020:295–336. https://doi.org/10.1016/B978-0-12-818438-7.00012-5.

Zhang R, Lu Y, Li Q. Clinical implementation challenges of large language models. npj Digit Med. 2024;7:144.

Sendak MP, D’Arcy J, Kashyap S, Gao M, Nichols M, Corey K, et al. A path for translation of machine learning products into healthcare delivery. EMJ Innov. 2020;4(1):48-55.

Si Y, Wang J, Zhao X. Evaluating safety of conversational AI in chronic disease management. npj Digit Med. 2024;7:98.

Labrique AB, Vasudevan L, Mehl G, Rosskam E, Hyder AA. Digital health and AI governance priorities. Lancet. 2018;392:1219-20.

Ning Y, Teixayavong S, Shang Y, Savulescu J, Nagaraj V, Miao D, et al. Generative artificial intelligence and ethical considerations in health care: a scoping review and ethics checklist. Lancet Digit Health. 2024;6(11):e848-e856.

Sheikh A, Anderson M, Albala S, Casadei B, Franklin BD, Richards M, et al. Health information technology and digital innovation for national learning health and care systems. Lancet Digit Health. 2021;3(6):e383-e396. https://doi.org/10.1016/S2589-7500(21)00005-4.

Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. https://doi.org/10.1038/s41591-018-0300-7.

Panteli D, Aissaoui N, et al. AI in public health policy. Lancet Public Health. 2024;9:e240-e250.

Chino Y, Tenenbaum J, Glusman G, et al. AI in medicine: creating a safe future. Lancet. 2023;402(10398):503-4.

Armoundas AA, Faghri PD. AI policy frameworks in healthcare. eClinicalMedicine. 2024;68:102312.

Mello MM, Cohen IG. Regulating AI in health care. JAMA. 2024;331(7):612-3.

Murray SG, Mishuris RG. Generative AI in clinical workflows. JAMA Netw Open. 2024;7:e241122.

Everson J, Lee SYD, Adler-Milstein J. Adoption of AI in US hospitals. Health Aff (Millwood). 2019;38(7):1133-40.

Nguyen TD, Whaley CM, Mehta N, et al. AI adoption trends in healthcare systems. JAMA Health Forum. 2024;5:e240091.

Wachter RM, Brynjolfsson E. Generative AI in health care. JAMA. 2024;331:65-6.

Zaretsky J, Kim JM, Baskharoun S, Zhao Y, Austrian J, Aphinyanaphongs Y, et al. Generative Artificial Intelligence to Transform Inpatient Discharge Summaries to Patient-Friendly Language and Format. JAMA Netw Open. 2024;7(3):e240357. https://doi.org/10.1001/jamanetworkopen.2024.0357.

Angus DC, Khera R, Lieu T, Liu V, Ahmad FS, Anderson B,et al. AI, Health, and Health Care Today and Tomorrow: The JAMA Summit Report on Artificial Intelligence. JAMA. 2025;334(18):1650-64. https://doi.org/10.1001/jama.2025.18490.