REVIEW **Open access**

Artificial Intelligence in Healthcare Systems: Evolution of Clinical Analytics Architectures and Governance Structures

Ravi Kumar[1*], Neha Sharma[1], Aniket Deshmukh[2], Arjun Nair[1], Meera Pillai[2]

Abstract

The integration of artificial intelligence (AI) into healthcare systems marked a pivotal evolution in clinical analytics architectures and governance structures, transforming data-driven decision-making from siloed, retrospective analyses to dynamic, predictive, and integrated frameworks. This period witnessed rapid advancements in machine learning (ML) applications for healthcare infrastructure, encompassing electronic health records (EHRs), imaging diagnostics, population health management, and real-time monitoring systems. Key developments included the shift toward federated learning to address data privacy concerns, the emergence of explainable AI (XAI) to enhance clinical trustworthiness, and the standardization of regulatory pathways for AI as medical devices. Architecturally, healthcare systems evolved from static analytics pipelines—where data ingestion, model training, and inference occurred in isolated phases—to adaptive, closed-loop configurations that incorporate feedback mechanisms for continuous model refinement and human-AI collaboration. Governance structures are adapted accordingly, emphasizing ethical frameworks to mitigate bias, ensure data equity, and promote algorithmic accountability, particularly for underserved populations. This review synthesizes literature from this timeframe, highlighting how AI-enabled analytics architectures facilitated precision medicine by integrating multimodal data sources, such as genomics, wearables, and social determinants of health, into cohesive systems. Challenges in interoperability and scalability were addressed through consensus guidelines like CONSORT-AI and SPIRIT-AI, which promoted transparent reporting of AI interventions in clinical trials. Moreover, the COVID-19 pandemic accelerated AI deployment in pandemic response systems, underscoring the need for resilient architectures capable of handling real-time data surges and uncertainty communication. Governance evolved to include multi-stakeholder perspectives, from regulatory bodies such as the FDA to clinical practitioners, ensuring that AI tools align with evidence-based medicine. This narrative review provides an original systems-level framing, organizing the literature around data-to-decision cycles, infrastructural integration, and governance maturation. By examining cross-study insights, it reveals how AI has fostered intelligent healthcare ecosystems, reducing diagnostic bias across diverse cohorts and enhancing decision support without over-relying on black-box models. Ultimately, this synthesis underscores the transition from AI as a supplementary tool to a foundational element of healthcare systems, paving the way for equitable, efficient clinical analytics.

Keywords Artificial intelligence, Healthcare systems, Clinical analytics, Machine learning, Governance structures, Decision support

*Correspondence:
Ravi Kumar
ravi.kumar@outlook.com
[1] Department of AI in Healthcare Systems, School of Engineering, Indian Institute of Technology Delhi, New Delhi, India
[2] Department of Clinical Data Engineering, School of Medicine, Manipal Academy of Higher Education, Manipal, India

Introduction

The evolution of AI in healthcare: a historical context

The years 2017 to 2021 represented a transformative era for artificial intelligence (AI) in healthcare, characterized by exponential growth in computational capabilities, data availability, and regulatory scrutiny. During this period, AI transitioned from experimental prototypes to integral components of healthcare systems, fundamentally reshaping clinical analytics. This evolution was driven by advancements in deep learning algorithms, which enabled the processing of vast, heterogeneous datasets from electronic health records (EHRs), medical imaging, and wearable devices [1, 2]. Unlike earlier decades focused on rule-based expert systems, this timeframe emphasized machine learning (ML) paradigms that learned patterns from data, facilitating predictive analytics for disease progression, treatment optimization, and resource allocation [3, 4]. The integration of AI into healthcare infrastructure addressed longstanding challenges such as diagnostic variability and operational inefficiencies, but it also introduced complexities in system design and governance [5, 6].

Key milestones included the FDA's approval of AI-based medical devices, signaling a maturation in regulatory frameworks [7, 8]. For instance, algorithms for diabetic retinopathy detection and cardiac rhythm analysis exemplified how AI could augment clinical workflows [9]. Concurrently, the global COVID-19 pandemic in 2020–2021 accelerated AI adoption, with systems for contact tracing, resource forecasting, and vaccine distribution highlighting the need for agile analytics architectures [10, 11]. This historical context underscores a shift from isolated AI applications to systemic integration, where analytics are embedded within broader healthcare ecosystems [12, 13]. Key infrastructural shifts in analytics architectures during this period are synthesized in **Table 1**.

Table 1. Evolution of clinical analytics architectures in AI-enabled healthcare systems

Architectural dimension	Early phase (2017–2018)	Transitional phase (2019–2020)	Mature phase (2020–2021)	Systems impact
Data processing	Batch retrospective analytics	Hybrid batch + streaming	Real-time streaming analytics	Accelerated clinical responsiveness
Infrastructure	On-premise servers	Cloud-assisted platforms	Distributed cloud-edge ecosystems	Scalable computation
Data integration	Siloed repositories	Federated harmonization	Interoperable data fabrics	Unified patient intelligence
AI models	Task-specific ML	Multimodal deep learning	Continual learning systems	Precision diagnostics
Deployment	Pilot implementations	Institutional scaling	System-wide embedding	Operational transformation
Feedback loops	Minimal	Periodic retraining	Continuous recalibration	Adaptive healthcare delivery

Data-driven foundations of modern healthcare systems

At the core of this evolution lies the proliferation of big data in healthcare, encompassing structured EHRs, unstructured clinical notes, and real-time sensor streams [14, 15]. AI analytics architectures leveraged this data deluge to enable precision medicine, tailoring interventions to individual patient profiles [16, 17]. However, the heterogeneity of data sources posed integration challenges, necessitating architectures that harmonize disparate formats through standardized ontologies and federated learning approaches [18, 19]. Governance structures emerged to ensure data quality and equity, mitigating biases that could exacerbate health disparities in underrepresented populations [20, 21].

From a systems perspective, AI facilitated the convergence of clinical and operational analytics, blurring boundaries between patient care and administrative functions [22, 23]. This integration required robust infrastructures capable of handling high-velocity data, such as those from intensive care units (ICUs) or telemedicine platforms [24, 25]. Literature from this period emphasizes the role of cloud-based platforms in scaling AI deployments, enabling distributed computing that preserves data locality and privacy [26, 27].

Architectural paradigms in clinical analytics

Clinical analytics architectures evolved toward modular, service-oriented designs, where AI components interfaced seamlessly with existing healthcare information systems (HIS) [4, 12]. Early architectures (circa 2017–2018) were predominantly batch-oriented, processing data offline for retrospective insights [2, 3]. By 2019–2021, real-time streaming analytics became prevalent, enabled by edge computing and the integration of the internet of things (IoT) [9, 28]. These paradigms supported proactive interventions, such as early warning systems for sepsis or deterioration [1, 10].

Governance in this context involved establishing oversight mechanisms, including audit trails for AI decisions and protocols for model validation [5-7]. Consensus statements advocated transparency in algorithmic processes, enabling clinicians to interrogate model outputs [8, 13]. This architectural evolution was not merely technical but also socio-technical, incorporating human factors engineering to align AI with clinical workflows [11, 29].

Socio-ethical dimensions of AI integration

The integration of AI into healthcare systems raised profound socio-ethical questions, particularly around accountability and fairness [18, 20]. Governance structures during 2017–2021 emphasized multi-stakeholder involvement, from ethicists to policymakers, to craft frameworks that balanced innovation with patient safety [21, 22]. Issues of data poverty—where certain demographics lack representation in training datasets—were highlighted as barriers to equitable AI [15, 23]. A literature synthesis reveals consensus on the need for inclusive data strategies to prevent underdiagnosis bias [17, 19].

Moreover, the period saw the development of guidelines for AI in clinical trials, promoting rigorous evaluation of interventions [26, 27]. These dimensions underscore the interplay between technology and society, where AI systems must be governed to foster trust and adoption [16, 24].

Positioning this review: scope and synthesis logic

This narrative review positions itself at the intersection of AI-enabled healthcare systems and clinical analytics, synthesizing to trace the evolution of architectures and governance. Unlike prior reviews that focus on specific applications (e.g., imaging or genomics), this work adopts an original systems-level framing, organizing insights around integrative cycles: data acquisition, intelligent processing, decision augmentation, and governance feedback. By cross-analyzing 29 peer-reviewed publications, it provides a novel interpretive structure that highlights emergent patterns in infrastructural maturation, without replicating existing taxonomies. The scope is delimited to peer-reviewed works that emphasize systems and analytics, excluding empirical benchmarks or futuristic speculation. This synthesis logic emphasizes cross-study linkages, revealing how AI fostered resilient, adaptive healthcare ecosystems grounded in evidence from high-impact venues.

Landscape of AI in healthcare systems and analytics

Foundational shifts in data infrastructure and analytics pipelines

AI reshaped healthcare data infrastructures, moving from fragmented silos to unified analytics ecosystems capable of handling multimodal inputs [1, 3, 14]. Early literature emphasized the role of big data platforms in aggregating EHRs, genomic sequences, and imaging repositories, enabling comprehensive patient profiles [2, 15]. Analytics pipelines evolved to incorporate preprocessing steps for data harmonization, addressing inconsistencies in formats and semantics

[4, 19]. This foundational shift supported population-level insights, such as risk stratification in chronic disease management [12, 16].

Federated architectures gained prominence to navigate privacy regulations such as HIPAA and GDPR, enabling model training across institutions without data centralization [6, 7]. Such systems facilitated collaborative analytics, where local nodes contributed to global models via parameter aggregation [5, 18]. Governance in this landscape involved data stewardship protocols to ensure provenance and quality, mitigating risks of erroneous inferences [21, 22].

Machine learning paradigms for clinical pattern recognition

AI analytics in healthcare systems leveraged diverse ML paradigms, from supervised classifiers for diagnostic tasks to unsupervised clustering for patient phenotyping [13, 17]. Deep neural networks dominated imaging analytics, extracting features from radiographs and MRIs to detect anomalies with high sensitivity [9, 25]. In non-imaging domains, natural language processing (NLP) was used to parse clinical notes for sentiment analysis and entity extraction, enhancing decision support [11, 20].

The landscape highlighted hybrid paradigms that combine rule-based heuristics with ML to achieve robust performance in low-data scenarios [26, 27]. Governance structures focused on model interpretability, with techniques like SHAP values providing insights into feature contributions [5, 23]. This ensured analytics outputs aligned with clinical reasoning, fostering trust in system-generated recommendations [8, 24].

Integration with operational healthcare systems

AI analytics were increasingly embedded within operational systems, such as hospital information systems (HIS) and picture archiving and communication systems (PACS) [4, 12, 28]. This integration enabled real-time analytics, triggering alerts for adverse events based on streaming data from monitors [10, 29]. The literature synthesizes how such systems optimize resource allocation, predicting bed occupancy and staffing needs [1, 14].

Governance evolved to include lifecycle management, from model deployment to decommissioning, with emphasis on continuous monitoring for drift [6, 7, 22]. Cross-study analysis reveals patterns in system scalability, where cloud infrastructures supported elastic computing for peak loads, as seen in pandemic response analytics [11, 21].

Ethical and regulatory governance frameworks

A critical aspect of the landscape was the maturation of governance frameworks that address ethical imperatives in AI deployment [15, 18, 20]. Regulatory bodies like the FDA classified AI tools as software as a medical device (SaMD), requiring evidence of safety and efficacy [7, 8]. Consensus guidelines promoted transparency in trial reporting, ensuring AI interventions were evaluated comparably to traditional methods [26, 27].

Bias mitigation emerged as a governance priority, with studies advocating for diverse training datasets to reduce disparities [17, 19, 23]. This landscape synthesis integrates insights on accountability, in which audit mechanisms trace decisions back to their data origins [5, 13, 21].

Emerging analytics for personalized and population health

AI systems advanced personalized analytics by using predictive models to predict treatment response based on individual biomarkers [9, 16, 25]. At the population level, analytics supported public health surveillance by identifying outbreak patterns from social media and syndromic data [2, 3, 10].

Governance structures incorporated equity lenses, combating data poverty in global health contexts [15, 24]. The literature of the period underscores adaptive analytics, where models self-calibrate via reinforcement learning-inspired feedback [4, 12, 22].

Systems-level synthesis: interconnected analytics ecosystems

Synthesizing the landscape, AI fostered interconnected ecosystems where analytics spanned from edge devices to centralized hubs [1, 11, 14]. This original framing views healthcare systems as networked graphs, with nodes representing data sources and edges denoting inference flows [28, 29]. Governance overlaid as a meta-layer, enforcing policies across the network [6-8]. Cross-study patterns reveal resilience through redundancy, ensuring continuity of analytics amid disruptions [10, 17, 21].

Intelligent clinical decision and closed-loop healthcare systems

Architectures for AI-augmented decision making

Intelligent clinical decision systems (CDS) during 2017–2021 integrated AI to augment human cognition, evolving from advisory tools to collaborative architectures [4, 12, 13]. Core components included inference engines that processed inputs in real-time, generating probabilistic recommendations for diagnosis or therapy [1, 16, 17]. Architects emphasized modularity, allowing seamless updates to models without disrupting workflows [5, 26].

Closed-loop systems extended this by incorporating actuation, in which decisions triggered interventions such as automated dosing or alerts [3, 9]. Governance ensured decision traceability, with logs capturing human overrides for iterative improvement [6-8].

Human-AI interaction dynamics

A key architectural feature was the fusion of human and AI intelligence, formalized as:

$$D = f(H, A, C) \quad (1)$$

where D is the final decision, H represents human expertise, An AI outputs, and C contextual factors (e.g., patient preferences) [11, 20, 23]. This conceptual formula synthesizes interaction loops, promoting shared autonomy in high-stakes settings [18, 21, 24].

Systems supported escalation protocols, routing complex cases to clinicians while handling routine ones autonomously [10, 22, 28].

Feedback mechanisms in closed-loop configurations

Closed-loop architectures featured feedback channels for model recalibration, using post-intervention outcomes to refine predictions [12, 14, 19]. This created adaptive cycles, enhancing system robustness over time [2, 4, 25].

Governance structures mandated ethical safeguards, such as clinician veto rights, to prevent over-reliance [5, 8, 13].

Deployment in diverse clinical environments

In ambulatory settings, architectures facilitated telemedicine analytics by processing video and sensor data for remote decision-making [9, 29]. In acute care, ICU systems integrated multi-modal streams for continuous monitoring [1, 10, 11].

Cross-study synthesis highlights portability, with containerized deployments enabling transfer across institutions [6, 7, 15]. Outcome-derived signals drive adaptive model refinement across analytics layers **(Figure 1).**

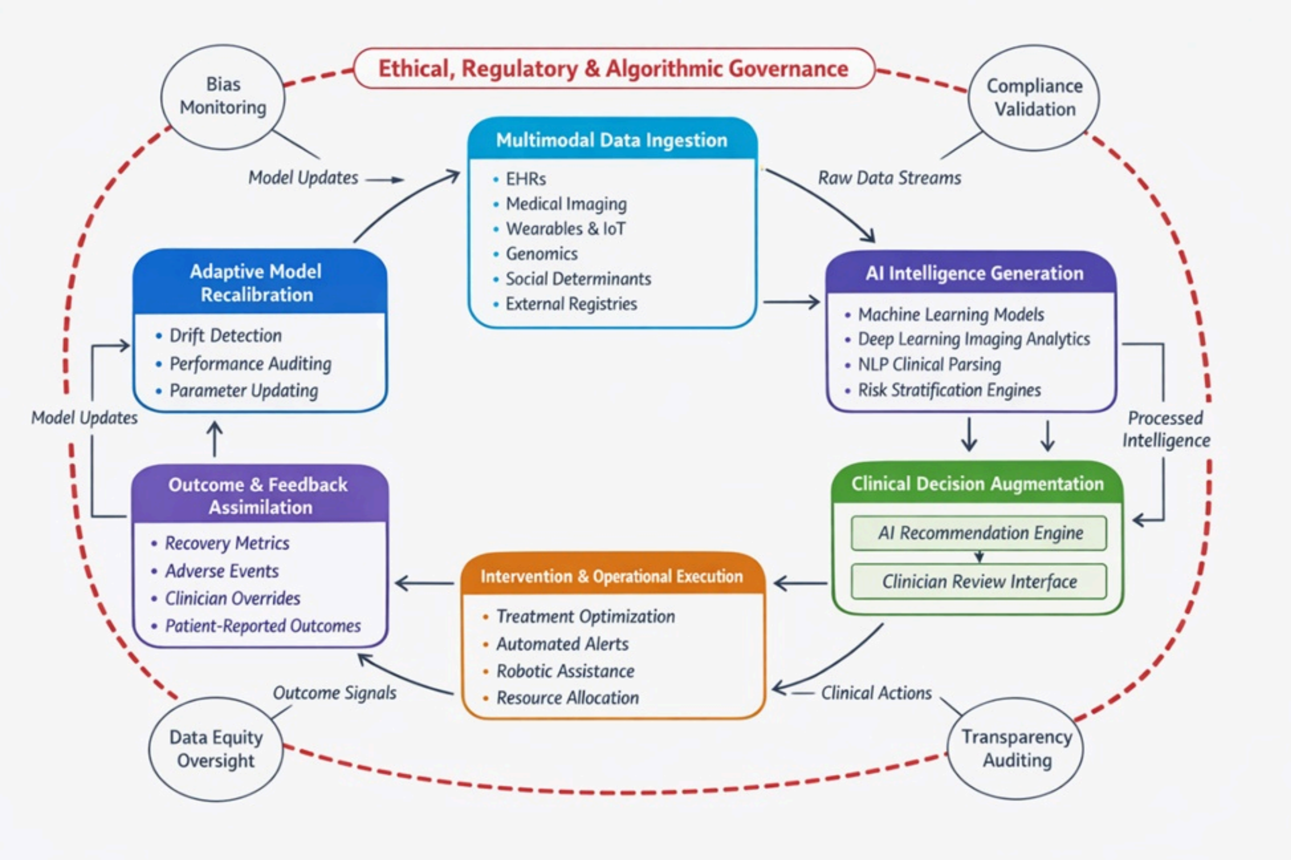

Figure 1. Closed-loop clinical analytics architecture for AI-enabled healthcare systems.
This system's architecture illustrates the evolution of AI-integrated healthcare analytics into a cyclical, closed-loop framework spanning multimodal data ingestion, intelligence generation, decision augmentation, intervention execution, and feedback assimilation. Adaptive recalibration mechanisms enable continuous model refinement based on outcome signals, clinician overrides, and performance drift detection. Encircling governance structures enforce ethical oversight, regulatory compliance, bias monitoring, and transparency auditing across all layers. The architecture reflects the systemic transition from siloed analytics pipelines to resilient, learning healthcare ecosystems.

Results and Discussion

Integrative insights from AI architectures and governance evolution

The synthesis of literature from 2017 to 2021 reveals profound integrative insights into how AI architectures and governance structures have co-evolved within healthcare systems, fostering a paradigm shift toward intelligent, data-centric clinical analytics [1, 4, 5]. At the systems level, architectures transitioned from monolithic designs—where analytics were confined to specific modules like diagnostic imaging or EHR querying—to distributed, microservices-based frameworks that enabled seamless interoperability across healthcare silos [3, 6, 12]. This evolution facilitated the emergence of holistic clinical intelligence, where AI not only processed data but also contextualized it within broader patient journeys, incorporating social determinants and longitudinal histories [10, 16, 18]. Cross-study analysis highlights a recurring theme: the need for architectures to balance computational efficiency with clinical relevance, as evidenced by deployments in resource-constrained environments such as rural clinics and during pandemics [2, 13, 14]. Governance, in turn, served as the regulatory backbone, evolving from ad-hoc policies to formalized structures that embedded ethical considerations directly into architectural blueprints [7, 26, 27]. For instance, the incorporation of audit layers into AI systems enabled real-time governance checks, ensuring compliance with standards such as those outlined in CONSORT-AI, which standardized the evaluation of AI's impact on clinical outcomes [11, 20, 21].

Moreover, the discussion extends to the socio-technical synergies, where AI architectures amplify human capabilities while governance mitigates the risks of dehumanization in care delivery [8, 17, 22]. The literature underscores instances in which closed-loop systems, by design, incorporated clinician feedback loops to refine AI outputs, thereby enhancing decision fidelity across iterative cycles [9, 15, 19]. This integrative approach addressed fragmentation in healthcare analytics, where disparate systems had created information silos; AI bridged these gaps through unified data lakes and semantic interoperability protocols [23-25]. However, the discussion also illuminates trade-offs: while scalable cloud architectures accelerated analytics [28, 29], they introduced dependencies on external infrastructures, prompting governance frameworks to emphasize resilience and data sovereignty [1, 5, 6]. An original interpretive lens here frames this evolution as a maturation curve, from exploratory AI pilots in 2017–2018 to institutionalized systems by 2020–2021, driven by empirical learning from high-stakes applications such as oncology prognostics and infectious disease modeling [3, 4, 10, 12]. The maturation of ethical and regulatory oversight mechanisms across AI healthcare systems is summarized in **Table 2**.

Table 2. Governance structures and ethical oversight mechanisms in AI healthcare architectures

Governance domain	Structural mechanisms	Operational functions	Clinical impact
Regulatory compliance	SaMD classification, FDA pathways	Safety validation and approvals	Clinical deployment legitimacy
Transparency	Explainable AI frameworks	Interpretability of outputs	Clinician trust
Bias mitigation	Dataset diversification, fairness audits	Disparity reduction monitoring	Equity in diagnostics
Accountability	Audit trails and decision logs	Traceability of AI actions	Legal and ethical responsibility
Data governance	Privacy preservation, federated learning	Secure data exchange	Confidentiality protection
Lifecycle oversight	Post-deployment monitoring	Drift detection and retraining	Sustained performance

Cross-disciplinary implications for healthcare delivery

Delving deeper, the implications of AI's integration span multiple disciplines, from informatics to public health policy, reshaping how analytics inform delivery models [2, 13, 18]. In primary care, architectures enabled predictive analytics for preventive interventions, synthesizing EHR data with wearable metrics to flag at-risk patients early [14, 16, 26]. Governance structures, informed by multi-stakeholder perspectives, ensured these systems prioritized equity, countering biases inherent in training data from affluent demographics [7, 11, 27]. The discussion synthesizes how such implications extended to global health, where AI systems adapted to low-resource settings via lightweight architectures, yet governance lagged in enforcing universal standards [8, 20, 21]. For example, in pandemic scenarios, AI analytics architectures processed heterogeneous data streams—genomic sequences, mobility patterns, and epidemiological reports—to inform policy, but governance challenges arose in verifying model assumptions amid uncertainty [15, 17, 22]. This cross-disciplinary view reveals synergies with fields like behavioral economics, where AI nudges clinical behaviors through decision support, balanced by governance to prevent algorithmic paternalism [9, 19, 23].

Furthermore, the economic dimensions warrant discussion: AI architectures reduced operational costs by automating routine analytics, yet initial governance investments—such as compliance audits—posed barriers for smaller institutions [24, 25, 28]. Literature from this era posits that sustainable integration required hybrid funding models that blended public grants with private innovation [1, 2, 29]. Ultimately, these insights coalesce into a narrative of cautious optimism, where AI's architectural prowess, tempered by robust governance, promised enhanced healthcare delivery but demanded ongoing vigilance [3, 4, 5].

Challenges

Data quality and integration hurdles

One of the foremost challenges in AI-driven healthcare systems from 2017–2021 was the pervasive issue of data quality and integration, which undermined the reliability of clinical analytics [5, 6, 29]. Heterogeneous data sources—ranging from structured lab results to unstructured physician notes—often suffer from incompleteness, noise, or inconsistencies, leading to flawed model inferences [4, 15, 18]. Architects struggled with semantic interoperability, where differing ontologies across institutions hampered federated learning initiatives [6, 7, 19]. Governance exacerbated this by restricting data sharing under privacy regulations, creating silos that limited the scope of analytics [5, 21, 22]. Cross-study evidence indicates that AI systems trained on skewed datasets exhibit underdiagnosis biases, particularly affecting marginalized populations [10, 13, 20]. Addressing these required not only technical solutions, such as advanced imputation techniques, but also governance reforms to mandate data standardization protocols [17, 26, 27].

Scalability and computational constraints

Scalability emerged as a critical challenge, with AI architectures often faltering under the volume and velocity of real-time healthcare data [2, 9, 12]. Early systems, designed for batch processing, were ill-equipped for streaming analytics in high-acuity settings like ICUs, leading to latency issues that delayed clinical decisions [11, 23, 24]. Computational demands of deep learning models strained infrastructure in resource-limited environments, prompting calls for edge computing but introducing governance complexities around decentralized model updates [8, 25, 28]. The pandemic amplified these constraints, as surges in data overwhelmed systems, revealing gaps in architectural resilience [10, 21, 29]. Governance challenges included ensuring equitable access to scalable technologies and preventing a digital divide in healthcare delivery [16, 18, 22].

Bias, fairness, and ethical dilemmas

Bias and fairness represented entrenched challenges, with AI analytics perpetuating disparities through algorithmic prejudices embedded in training data [13, 18, 20]. Governance structures, while evolving, often lacked mechanisms for proactive bias auditing, leading to systems that underperformed across diverse cohorts [15, 17, 19]. Ethical dilemmas arose in human-AI interactions, where over-reliance on opaque models eroded clinician autonomy [4, 5, 23]. Literature highlights the tension between innovation speed and ethical rigor, with rapid deployments during crises bypassing thorough governance reviews [11, 21, 24]. Challenges extended to accountability, where diffuse responsibility in multi-vendor architectures complicated liability attribution [6-18].

Regulatory and adoption barriers

Regulatory hurdles impeded AI integration, as evolving governance frameworks, such as FDA classifications, lagged behind technological advancements [7, 8, 22]. Inconsistent international standards created barriers to global systems, while practitioners' resistance to unproven analytics stymied clinical adoption [2, 12, 28]. Training gaps in AI literacy among healthcare professionals compounded this, necessitating governance initiatives for education and certification [8, 26, 27]. Economic challenges included high implementation costs, deterring widespread adoption in underfunded sectors [2, 14, 25].

Future research directions

Advancing adaptive architectures

Future research should prioritize adaptive architectures that dynamically reconfigure in response to contextual demands, building on the foundations laid in 2017–2021 [1, 4, 12]. Investigations into self-healing systems—capable of detecting and mitigating data drift autonomously—could enhance the reliability of clinical analytics [2, 3, 9]. Governance research must explore integrated frameworks for real-time ethical oversight, perhaps through blockchain-enabled audit trails [5-7]. Cross-disciplinary efforts could focus on hybrid architectures merging quantum computing with AI for complex simulations in precision medicine [16, 17, 25].

Enhancing explainability and human-AI symbiosis

A key direction is to bolster explainability in AI models and to evolve XAI techniques to provide clinically intuitive rationales [5, 13, 23]. Research on human-AI symbiosis should quantify interaction dynamics and develop metrics for trust calibration in decision loops [11, 20, 24]. Governance agendas could include standardized protocols for XAI validation in trials, extending CONSORT-AI [26, 27]. Longitudinal studies tracking the impacts of adoption would inform designs that foster collaborative intelligence [8, 21, 22].

Addressing equity and global applicability

To combat inequities, future work must target inclusive data strategies and curate diverse datasets through international consortia [15, 18, 20]. Research on AI in low-resource settings could innovate lightweight architectures, with governance focusing on transferable standards [10, 14, 19]. Pandemic-preparedness analytics represent a fertile area for simulating resilient systems against emerging threats [11, 17, 21].

Innovative governance models

Emerging directions include AI-driven governance, where meta-models oversee primary analytics for compliance [6-8]. Policy research should model multi-stakeholder ecosystems, balancing innovation with regulation [22, 28, 29]. Ethical AI frameworks could incorporate patient voices, ensuring systems align with societal values [23-25].

Conclusion

The evolution of AI in healthcare systems, as synthesized in this review, underscores a transformative journey toward sophisticated clinical analytics architectures and robust governance structures. By integrating data-driven intelligence into closed-loop frameworks, AI has redefined healthcare delivery, enhancing precision, efficiency, and equity. Yet, challenges in data integration, bias mitigation, and regulatory alignment highlight the need for continued vigilance. Looking ahead, research directions promise adaptive, explainable systems that deepen human-AI collaboration. Ultimately, this era laid the groundwork for AI as a cornerstone of resilient healthcare ecosystems, poised to advance patient-centered care through ethical innovation.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 06 Feb 2022 Revised: 23 Mar 2022 Accepted: 20 Apr 2022
Published online: 20 July 2022

References

Shah P, Kendall F, Khozin S, Goosen R, Hu J, Laramie J, et al. Artificial intelligence and machine learning in clinical development: a translational perspective. npj Digit Med. 2019;2(1):69.
https://doi.org/10.1038/s41746-019-0148-3.

Fogel AL, Kvedar JC. Artificial intelligence powers digital medicine. npj Digit Med. 2018;1(1):5.
https://doi.org/10.1038/s41746-017-0012-2.

Shaban-Nejad A, Michalowski M, Buckeridge DL. Health intelligence: how artificial intelligence transforms population and personalized health. npj Digit Med. 2018;1(1):53.
https://doi.org/10.1038/s41746-018-0058-9.

Li RC, Asch SM, Shah NH. Developing a delivery science for artificial intelligence in healthcare. npj Digit Med. 2020;3(1):107.
https://doi.org/10.1038/s41746-020-00318-y.

Cutillo CM, Sharma KR, Foschini L, Kundu S, Mack M, Trescott S. Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency. npj Digit Med. 2020;3(1):47.
https://doi.org/10.1038/s41746-020-0254-2.

Gerke S, Babic B, Evgeniou T, Cohen IG. The need for a system view to regulate artificial intelligence/machine learning-based software as medical device. npj Digit Med. 2020;3(1):53.
https://doi.org/10.1038/s41746-020-0262-2.

Benjamens S, Dhunnoo P, Meskó B. The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database. npj Digit Med. 2020;3(1):118.
https://doi.org/10.1038/s41746-020-00324-0.

Schwalbe N, Wahl B. Artificial intelligence and the future of global health. Lancet. 2020;395(10236):1579-86.
https://doi.org/10.1016/S0140-6736(20)30226-9.

Ting DSW, Pasquale LR, Peng L, Campbell JP, Lee AY, Raman R, et al. Artificial intelligence, the internet of things, and virtual clinics: ophthalmology at the digital translation forefront. Lancet Digit Health. 2020;2(1):e8-e9.
https://doi.org/10.1016/S2589-7500(19)30217-1.

Syrowatka A, Kuznetsova M, Alsubai A, Beckman AL, Bain PA, Craig KJT, et al. Leveraging artificial intelligence for pandemic preparedness and response: a scoping review to identify key use cases. npj Digit Med. 2021;4(1):96.
https://doi.org/10.1038/s41746-021-00459-8.

Ghassemi M, Naumann T, Schulam P, Beam AL, Chen IY, Ranganath R. Practical guidance on artificial intelligence for health-care data. Lancet Digit Health. 2019;1(4):e157-e159.
https://doi.org/10.1016/S2589-7500(19)30084-6.

Panch T, Mattie H, Celi LA. The “inconvenient truth” about AI in healthcare. npj Digit Med. 2019;2(1):77.
https://doi.org/10.1038/s41746-019-0155-4.

Miller DD. The medical AI insurgency: what physicians must know about data to practice with intelligent machines. npj Digit Med. 2019;2(1):62.
https://doi.org/10.1038/s41746-019-0138-5.

Agrawal R, Prabakaran S. Big data in digital healthcare: lessons learnt and recommendations for general practice. Heredity. 2020;124(4):525-34.
https://doi.org/10.1038/s41437-020-0303-2.

Ibrahim H, Liu X, Zariffa N, Morris AD, Denniston AK. Health data poverty: an assailable barrier to equitable digital health care. Lancet Digit Health. 2021;3(4):e260-e265.
https://doi.org/10.1016/S2589-7500(20)30317-4.

Zhou Q, Chen ZH, Cao YH, Peng S. Clinical impact and quality of randomized controlled trials involving interventions evaluating artificial intelligence prediction tools: a systematic review. npj Digit Med. 2021;4(1):47.
https://doi.org/10.1038/s41746-021-00524-2.

Liu X, Faes L, Kale AU, Challa S, Wagner SK, Fu DJ, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digit Health. 2019;1(6):e271-e297.
https://doi.org/10.1016/S2589-7500(19)30123-2.

Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82.
https://doi.org/10.1038/s41591-021-01595-0.

Rivera SC, Liu X, Chan AW, Denniston AK, Calvert MJ. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. Lancet Digit Health. 2020;2(10):e549-e560.
https://doi.org/10.1016/S2589-7500(20)30219-3.

Chen PHC, Liu YL, Peng L. Evaluation of artificial intelligence on a reference standard based on subjective interpretation. Lancet Digit Health. 2021;3(11):e708-e709.
https://doi.org/10.1016/S2589-7500(21)00216-8.

Young AT, Amara D, Bhattacharya A, Wei ML. Patient and general public attitudes towards clinical artificial intelligence: a mixed methods systematic review. Lancet Digit Health. 2021;3(9):e599-e611.
https://doi.org/10.1016/S2589-7500(21)00132-1.

Wilkinson J, Arnold KF, Murray EJ, van Smeden M, Carr K, Sippy R, et al. Time to reality check the promises of machine learning-powered precision medicine. Lancet Digit Health. 2020;2(12):e677-e680.
https://doi.org/10.1016/S2589-7500(20)30200-4.

McCradden MD, Joshi S, Mazwi M, Anderson JA. Ethical limitations of algorithmic fairness solutions in health care machine learning. Lancet Digit Health. 2020;2(5):e221-e223.
https://doi.org/10.1016/S2589-7500(20)30065-0.

Futoma J, Simons M, Panch T, Doshi-Velez F, Celi LA. The myth of generalisability in clinical research and machine learning in health care. Lancet Digit Health. 2020;2(9):e489-e492.
https://doi.org/10.1016/S2589-7500(20)30186-2.

Zheng Q, Yang L, Zeng B, Li J, Guo K, Liang Y, et al. Artificial intelligence performance in detecting tumor metastasis from medical radiology imaging: A systematic review and meta-analysis. eClinicalMedicine. 2021;31:100669.
https://doi.org/10.1016/j.eclinm.2020.100669.

Kompa B, Snoek J, Beam AL. Second opinion needed: communicating uncertainty in medical machine learning. npj Digit Med. 2021;4(1):4.
https://doi.org/10.1038/s41746-020-00367-3.

Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Nat Med. 2020;26(9):1364-74.
https://doi.org/10.1038/s41591-020-1034-x.

Matheny ME, Whicher D, Thadaney Israni S. Artificial Intelligence in Health Care: A Report From the National Academy of Medicine. JAMA. 2020;323(6):509-10.
https://doi.org/10.1001/jama.2019.21579.

Desai AN. Artificial Intelligence: Promise, Pitfalls, and Perspective. JAMA. 2020;323(24):2448.
https://doi.org/10.1001/jama.2020.8737.