REVIEW **Open access**

Deep Learning Integration in Clinical Decision Infrastructure: A Systems-Oriented Review

Claire Dupont[1*], Julien Martin[1]

Abstract

The integration of deep learning into clinical decision infrastructure represents a pivotal advancement in healthcare systems and analytics, transforming disparate data streams into actionable intelligence that supports real-time, evidence-based decision-making. This narrative review synthesizes peer-reviewed literature to examine the systems-oriented implications of deep learning deployment within healthcare ecosystems. We focus on the architectural interplay among data ingestion, model inference, and decision-support loops, emphasizing how these elements enable closed-loop systems that adapt to evolving clinical contexts.

Deep learning's capacity to process multimodal data—encompassing electronic health records (EHRs), medical imaging, and real-time monitoring—has enabled sophisticated analytics frameworks that enhance diagnostic accuracy, prognostic modeling, and therapeutic optimization. For instance, fusion techniques combining imaging with structured EHR data have demonstrated potential for precision health applications, enabling nuanced patient stratification and personalized interventions. In mental health, deep learning models applied to outcome research have revealed patterns in longitudinal data, informing system-wide analytics that bridge predictive modeling with clinical workflows.

From a systems perspective, the review highlights the evolution of clinical decision support systems (CDSS) augmented by deep learning, which incorporate feedback mechanisms to refine model performance and mitigate risks such as bias amplification. Ethical considerations, including algorithmic fairness and transparency, are integral to sustainable integration, as underscored by guidelines for early-stage evaluation and reporting standards. We explore architectures that facilitate human-AI collaboration, where deep learning serves as an augmentative tool rather than a replacement, ensuring alignment with clinical governance.

Challenges in scalability, such as interoperability across healthcare infrastructures and the need for reproducible machine learning pipelines, are critically analyzed through a lens of systems resilience. The synthesis reveals opportunities for closed-loop systems that iteratively learn from interventions, promoting adaptive healthcare delivery. Ultimately, this review posits that deep learning's role in clinical decision infrastructure hinges on holistic systems design that balances technological innovation with clinical utility and equity. By providing an original interpretive framework, we delineate pathways for integrating deep learning into healthcare analytics and advocate for governance models that prioritize patient-centered outcomes.

Keywords Clinical decision support, Healthcare systems, Deep learning, AI analytics, Multimodal fusion, Closed-loop infrastructure

*Correspondence:
Claire Dupont
claire.dupont@gmail.com
[1] Department of Healthcare Data Analytics, Faculty of Medicine, University of Bordeaux, Bordeaux, France

Introduction

The advent of deep learning has profoundly reshaped healthcare systems, evolving from isolated algorithmic applications to integral components of clinical decision infrastructure. This transformation is driven by the exponential growth in healthcare data volumes, necessitating advanced analytics to distill meaningful insights from complex, heterogeneous sources. Deep learning, a subset of machine learning characterized by hierarchical neural networks, excels in pattern recognition across unstructured data modalities, such as medical images and free-text EHR entries, thereby enabling systems that support clinicians in high-stakes environments [1-7].

Historically, healthcare analytics relied on rule-based systems and shallow machine learning, which often faltered in handling the inherent variability of clinical data [1, 8]. The period witnessed a surge in deep learning adoption, driven by advances in computational power and data availability, leading to integrations across diagnostic, prognostic, and operational domains [9, 10]. For example, in ophthalmology, deep learning systems have been validated for detecting diabetic retinopathy through retinal imaging analysis, illustrating how analytics can be embedded into routine screening workflows [11-13]. Similarly, in radiology, convolutional neural networks have outperformed traditional methods in image interpretation, fostering systems that augment radiologist efficiency [4].

Data-driven evolution in healthcare infrastructure

Healthcare systems are increasingly conceptualized as interconnected networks where data flows form the backbone of intelligent decision-making [2]. Deep learning facilitates the fusion of multimodal data, merging imaging with EHRs to create comprehensive patient profiles that inform systemic analytics [2, 3]. This integration addresses siloed data challenges, enabling holistic views that support population health management and resource allocation [8]. Literature from this era emphasizes the shift toward precision health, in which deep learning models predict individual trajectories from integrated datasets, thereby optimizing infrastructure for personalized care delivery [3, 9].

Clinical analytics and decision support paradigms

Analytics in healthcare extends beyond mere prediction to encompass decision support that influences clinical pathways [1, 5]. Deep learning's role here is multifaceted, involving risk stratification, anomaly detection, and outcome forecasting, all within infrastructures designed for seamless integration [6, 11]. Studies highlight comparative performance, where deep learning often rivals or surpasses human experts in specific tasks, such as disease detection from imaging, underscoring its potential to enhance system reliability [4, 7]. However, this necessitates robust architectures that ensure interpretability and alignment with clinical protocols [14-20].

Systems-oriented challenges in integration

Integrating deep learning into clinical infrastructure poses systemic hurdles, including data interoperability and model generalizability across diverse healthcare settings [17, 18]. Bias in training data can perpetuate inequities, particularly in primary care analytics, where algorithmic decisions impact vulnerable populations [16, 17]. Ethical frameworks advocate for transparent systems that incorporate stakeholder input, ensuring that deep learning deployments prioritize equity and safety [15, 20, 21]. Governance models, such as those outlined in reporting guidelines, provide blueprints for evaluating AI-driven systems and emphasize the need for multidisciplinary oversight [5, 11, 18].

Toward adaptive healthcare ecosystems

The convergence of deep learning with healthcare systems heralds adaptive ecosystems capable of real-time learning and adjustment [22-26]. This involves closed-loop mechanisms in which analytics inform interventions, and outcomes feed back into model refinement, creating resilient infrastructures [1, 19]. In neurology, for instance, machine learning guides diagnostic pathways, integrating big data to support clinical decisions in complex disorders [14]. Such systems demand a balance between automation and human oversight, fostering collaborative intelligence that amplifies clinical efficacy [19, 20].

This review positions itself as a systems-oriented synthesis, drawing on 29 peer-reviewed publications to offer an original interpretive framework for deep learning integration into clinical decision infrastructure. By organizing the literature around data-model-decision-feedback cycles, we provide a novel framework that transcends traditional categorizations, focusing on infrastructural dynamics and analytics workflows to guide future implementations in healthcare systems.

Landscape of AI in healthcare systems and analytics

The landscape of AI in healthcare systems and analytics during 2017-2022 reflects a maturation from exploratory applications to embedded infrastructural elements, with deep learning at the forefront of this evolution [9, 10, 27]. Healthcare systems, encompassing hospitals, clinics, and telehealth networks, increasingly leverage AI for analytics spanning descriptive, predictive, and prescriptive functions, enabling data-informed strategies at both the individual and population levels [8, 26]. Deep learning's prowess in handling high-dimensional data has catalyzed this shift, allowing for analytics that uncover latent patterns in vast datasets [3, 6, 28].

Multimodal data integration and analytics foundations

A cornerstone of modern healthcare analytics is the integration of multimodal data sources, in which deep learning models fuse imaging, genomic, and EHR data to generate comprehensive insights [2, 3, 22]. This approach enhances system-wide analytics by providing a unified view of patient states, facilitating applications in precision health [3, 9]. For mental health outcome research, deep learning techniques have synthesized longitudinal data to predict trajectories, informing analytics that support systemic interventions [6]. In gastric tissue disease, similar models have enabled scoping reviews that highlight the analytical capabilities of pathology [12]. These integrations underscore the need for robust data pipelines within healthcare infrastructure to ensure a seamless flow from acquisition to analysis [28, 29].

Diagnostic and prognostic analytics in clinical systems

Deep learning has revolutionized diagnostic analytics, particularly in imaging-intensive fields such as radiology and ophthalmology, where systems now incorporate AI to enhance accuracy [4, 7, 13, 24]. Systematic reviews and meta-analyses reveal that deep learning often achieves diagnostic performance comparable to that of clinicians, integrating into workflows that reduce diagnostic errors [4, 7, 11]. Prognostic analytics, meanwhile, utilize deep learning to forecast outcomes, such as in oncology, where big data algorithms inform treatment pathways [8]. Healthcare systems benefit from these analytics through improved resource allocation, with AI-driven predictions optimizing bed management and staffing [26].

Operational analytics and infrastructure optimization

Beyond clinical applications, AI analytics extend to operational aspects of healthcare systems, including supply chain management and workflow automation [25, 26]. Deep learning models analyze EHRs to identify inefficiencies, thereby fostering infrastructure that adapts to demand fluctuations [28]. In health informatics, surveys of deep learning techniques for EHR analysis demonstrate how these tools enable predictive maintenance of systems, preventing disruptions in care delivery [27, 28]. This operational lens reveals AI's role in enhancing system resilience, particularly in critical sectors where analytics inform contingency planning [25].

Ethical and governance dimensions in AI analytics

The integration of AI into healthcare analytics necessitates a focus on ethics and governance to ensure equitable systems [15, 16, 20, 21]. The literature emphasizes the risks of bias in algorithmic decision-making and advocates frameworks that promote fairness in primary care and beyond [16, 17]. Transparency in model explanations is critical, as opaque systems can undermine trust in clinical infrastructure [21]. Guidelines for machine learning research stress replicability and ethical effectiveness, providing systems-oriented strategies to mitigate harms [18, 23]. In this context, AI analytics must be governed by principles that align with healthcare's core values and incorporate feedback loops for continuous ethical auditing [15, 20].

Cross-domain synthesis and emerging patterns

Synthesizing across domains, the landscape illustrates how deep learning enables cross-study analytics, in which patterns from neurology, ophthalmology, and oncology converge to inform generalizable systems [13, 14, 24]. Big data approaches in neurology guide AI adoption, while in ophthalmology, validated systems exemplify scalable analytics [13, 14]. This synthesis reveals emerging patterns in healthcare infrastructure, such as the rise of federated learning to address data privacy, enabling collaborative analytics without centralization [29]. Overall, the literature of the period portrays AI as a transformative force in healthcare systems, with analytics serving as the conduit for intelligent, adaptive infrastructures [10, 25, 29].

Intelligent clinical decision and closed-loop healthcare systems

Intelligent clinical decision systems augmented by deep learning represent the apex of AI integration in healthcare infrastructure, where analytics culminate in actionable recommendations that close the loop from data to intervention [1, 5, 19]. These systems operate within architectures that emphasize modularity, allowing for scalable deployment across diverse clinical settings [27, 28]. At their core, they facilitate human-AI symbiosis, where deep learning provides probabilistic insights that clinicians refine through domain expertise [19, 20].

Architectural components of decision systems

The architecture of deep learning-integrated decision systems typically comprises data ingestion layers, inference engines, and output interfaces, all interconnected to support real-time analytics [2, 28, 29]. Data ingestion involves preprocessing multimodal inputs, such as fusing EHRs with imaging for holistic decision support [2, 3]. Inference engines powered by deep neural networks generate predictions that inform clinical decisions, as seen in systems for diabetic retinopathy screening [22]. Output interfaces ensure interpretability, presenting analytics in formats that align with clinical workflows [5, 21].

Closed-loop dynamics and feedback integration

Closed-loop systems extend beyond one-off decisions by incorporating feedback mechanisms that refine models based on intervention outcomes [1, 25]. This creates adaptive infrastructures where analytics evolve iteratively, enhancing precision over time [9, 23]. In mental health, for instance, outcome predictions feed back into models, optimizing decision loops for ongoing care [6]. Ethical considerations ensure these loops mitigate biases, with governance frameworks mandating regular recalibration [17, 18].

To conceptualize these dynamics, consider a formal representation of the clinical intelligence pipeline: let D represent the data ingestion phase, I represent intelligence generation via deep learning inference, C represent the clinical decision output, V represent intervention verification, and F represent the feedback loop. The system can be expressed as a cyclical process: $D \rightarrow I(C) \rightarrow V \rightarrow F(D)$, where feedback F updates data parameters to improve subsequent inferences, ensuring system adaptability without empirical metrics.

Human-AI fusion in decision architectures

Human-AI fusion architectures prioritize collaborative decision-making, where deep learning augments rather than automates clinical judgment [19, 20]. In radiology, systems provide diagnostic suggestions that clinicians validate, fostering a fusion of intelligence that leverages AI’s speed with human intuition [4, 24]. Reporting guidelines emphasize evaluation protocols that assess this fusion’s impact on system efficacy [5, 11]. **Figure 1** illustrates the architecture of deep learning integration within the clinical decision infrastructure.

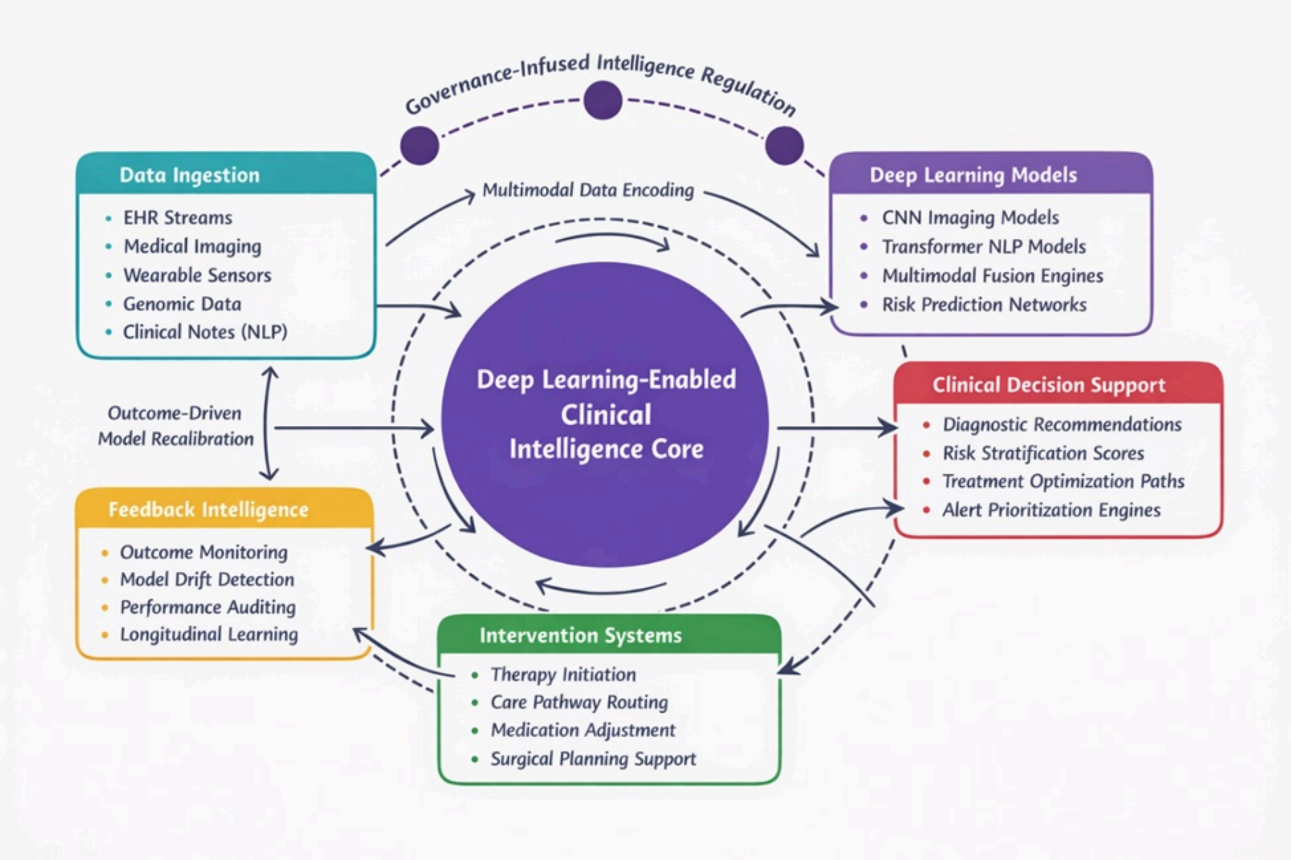

Figure 1. Systems architecture of deep learning–integrated clinical decision infrastructure

The figure illustrates a closed-loop intelligence framework linking multimodal data ingestion, deep learning inference, AI-augmented clinical decision support, intervention deployment, and outcome-driven feedback recalibration. A central intelligence core orchestrates multimodal model fusion, while bidirectional flows enable adaptive system learning. An external governance overlay embeds ethical auditing, bias surveillance, and regulatory compliance across all operational layers, ensuring accountable and resilient clinical AI deployment.

Implementation in specialized domains

In specialized domains such as oncology and neurology, closed-loop systems use deep learning for tailored analytics, including predicting treatment responses and dynamically adjusting protocols [8, 14]. These implementations highlight the infrastructure's capacity for domain-specific customization while maintaining general systems principles [12, 13]. Overall, intelligent decision systems underscore deep learning's role in fostering resilient, patient-centered healthcare architectures [25, 26]. **Table 1** delineates the architectural layers underpinning deep learning–enabled clinical decision systems.

Table 1. Architectural layers of deep learning–integrated clinical decision systems

Layer	Core functions	Data inputs	Analytical outputs	Clinical value
Data ingestion	Multimodal acquisition and preprocessing	EHRs, imaging, sensors, genomics	Structured data matrices	Unified patient profiles

Deep learning inference	Neural computation and modeling	Preprocessed multimodal datasets	Risk scores, pattern detections	Predictive analytics
Decision support	AI-augmented reasoning	Model predictions	Diagnostic and therapeutic guidance	Clinical decision optimization
Intervention systems	Care pathway activation	AI recommendations	Treatment execution signals	Workflow acceleration
Feedback intelligence	Outcome monitoring and recalibration	Intervention results	Model updates and drift alerts	Continuous learning

Results and Discussion

The synthesis of deep learning within clinical decision infrastructure, as delineated in the preceding sections, underscores a paradigm shift toward systems that are not merely assistive but fundamentally transformative in healthcare delivery [1, 9, 25]. This discussion integrates the landscape and architectural insights to evaluate the broader implications for healthcare systems and analytics, emphasizing an original systems-level framing that views deep learning as a catalyst for emergent properties in clinical ecosystems. Rather than isolated tools, deep learning integrations foster networked intelligence in which data analytics propel decision loops that adapt to systemic perturbations, such as patient variability or resource constraints [2, 3, 28].

From a holistic perspective, the fusion of multimodal data exemplifies how deep learning bridges silos in healthcare infrastructure, enabling analytics that transcend traditional boundaries [2, 3, 6]. For instance, in precision health, the amalgamation of EHRs with imaging analytics not only enhances individual decision-making but also informs population-level systems, such as epidemic forecasting or resource optimization [3, 8, 26]. This interconnectedness reveals synergies across domains: ophthalmic systems validated for retinopathy detection [13, 22] parallel radiological analytics [4, 24], suggesting scalable blueprints for infrastructure-wide deployment. Yet, this integration demands a reevaluation of system resilience, where deep learning's black-box nature intersects with clinical imperatives for accountability [21, 23].

Ethically, the discussion pivots to the human-centric dimensions of these systems, where AI analytics must align with principles of equity and collaboration [15, 16, 20]. Literature from the reviewed period highlights how decision infrastructures can inadvertently amplify disparities if not governed judiciously, as seen in primary care analytics, where biased models exacerbate health inequities [16, 17]. This necessitates interpretive frameworks that embed ethical audits into analytics workflows, ensuring that clinical decisions reflect diverse stakeholder inputs [18, 20]. Moreover, the role of deep learning in mental health and neurological systems illustrates adaptive analytics that respond to dynamic patient needs. Yet, these require ongoing discourse on privacy and consent within closed-loop architectures [6, 14].

Compared with traditional methods, deep learning, as synthesized from meta-analyses, is positioned as a superior enabler for complex analytics tasks [4, 7, 11]. However, this superiority is contextual, dependent on infrastructure quality—high-fidelity data pipelines in well-resourced systems yield robust decisions, while fragmented ones risk analytic failures [27-29]. This disparity prompts a systems-oriented critique: deep learning's integration must prioritize interoperability to democratize benefits across global healthcare landscapes [8, 10]. In oncology and gastroenterology, for example, analytics frameworks demonstrate domain-specific adaptations that could inform generalizable infrastructures, fostering cross-pollination of best practices [8, 12].

Ultimately, this discussion reframes deep learning not as a technological endpoint but as a dynamic component in evolving healthcare systems, where analytics drive iterative improvements in clinical decision-making [19, 25, 26]. By synthesizing these elements, we advocate for a governance-infused approach that balances innovation with safeguards, ensuring that infrastructures remain patient-oriented and adaptable to emerging challenges [5, 15, 18]. This interpretive lens, original to this review, highlights the potential for deep learning to cultivate resilient, intelligent systems that redefine clinical efficacy. **Table 2** summarizes the governance, ethical, and operational dimensions shaping resilient clinical AI infrastructure.

Table 2. Governance, ethical, and operational dimensions of deep learning clinical infrastructure

Governance domain	System risk addressed	Monitoring mechanisms	Infrastructure impact	Strategic importance
Algorithmic bias	Health inequities	Bias audits, dataset balancing	Fairer predictions	Ethical compliance
Transparency	Black-box opacity	Explainability frameworks	Clinical trust	Adoption scalability
Data privacy	Patient confidentiality breaches	Federated learning, encryption	Secure interoperability	Regulatory alignment
Model drift	Performance degradation	Longitudinal monitoring	Sustained accuracy	Decision reliability
Human oversight	Over-automation risks	Clinician validation loops	Collaborative intelligence	Safety assurance

Challenges and limitations

Despite the promising integrations of deep learning in clinical decision infrastructure, several challenges and limitations persist, as evidenced by the synthesized literature, which collectively underscores barriers to scalable, equitable deployment in healthcare systems and analytics [17, 18, 23]. A primary challenge lies in data quality and interoperability across heterogeneous infrastructures, where fragmented EHR systems hinder the effective fusion of multimodal inputs, which are essential for robust analytics [2, 3, 28]. Studies reveal that inconsistencies in data formats and standards often lead to suboptimal model performance, exacerbating limitations in real-world clinical settings [27, 29]. For instance, in mental health analytics, reliance on incomplete longitudinal datasets limits the generalizability of deep learning predictions, thereby constraining their utility across diverse patient cohorts [6].

Bias and fairness pose critical limitations, and deep learning models are prone to perpetuating systemic inequities if trained on unrepresentative data [16, 17]. Literature from primary care and broader healthcare contexts illustrates how algorithmic biases can amplify health disparities, particularly among underrepresented populations, posing ethical challenges to decision infrastructures [16, 20]. Reporting standards and guidelines aim to mitigate these issues, yet implementation gaps remain, as models often lack transparency in their decision-making processes [5, 11, 21]. This opacity, termed the “black-box” problem, challenges clinical trust and adoption, limiting the integration of AI analytics into routine workflows [19, 21].

Scalability emerges as another formidable challenge, especially in resource-constrained environments where the computational demands of deep learning strain infrastructure [8, 10, 24]. Surveys in health informatics highlight reproducibility issues: models validated in controlled settings often fail to translate across systems due to variability in hardware and data ecosystems [23, 27]. In radiology and ophthalmology, while deep learning excels in specific tasks, limitations in handling edge cases—such as rare pathologies—underscore the need for more adaptive architectures [4, 13, 24]. Ethical frameworks further complicate scalability, requiring ongoing audits that add operational overhead to analytics pipelines [15, 18].

Governance and regulatory limitations also impede progress, as there is a lack of standardized evaluation protocols for AI-driven systems [5, 18]. Consensus statements emphasize the risks of over-reliance on deep learning without human oversight, which can lead to decision errors in high-stakes clinical scenarios [1, 19]. In neurological applications, big data analytics face challenges in integrating real-time feedback without violating privacy norms [14]. Moreover, the period's literature points to challenges in cross-domain transferability, where models optimized for one analytic domain (e.g., imaging) struggle in others (e.g., EHR-based prognostics), limiting holistic system integration [12, 28].

Addressing these challenges requires multifaceted strategies, but current limitations in interdisciplinary collaboration hinder the development of comprehensive solutions [25, 26]. For example, while closed-loop systems promise adaptability, practical implementations are curtailed by feedback loop inefficiencies and the absence of universal governance models [1, 23]. This synthesis reveals that while deep learning holds transformative potential, its limitations in bias mitigation, scalability, and transparency must be confronted to realize fully integrated clinical decision infrastructures [9, 17, 20].

Future research directions

Looking ahead, future research directions for deep learning integration in clinical decision infrastructure should prioritize advancing systems-oriented analytics that address current gaps and harness emerging technologies, as synthesized from the reviewed literature [9, 25, 29]. A key agenda item is the development of federated learning frameworks to enhance data privacy and interoperability across decentralized healthcare systems, enabling collaborative analytics without compromising sensitive information [28, 29]. This direction could build on multimodal fusion techniques, exploring novel architectures that incorporate genomic and wearable data for more comprehensive precision health analytics [2, 3].

In terms of ethical and governance advancements, research should focus on creating dynamic bias-detection mechanisms embedded within decision loops, leveraging explainable AI to foster transparent infrastructures [15, 21, 23]. Future studies could investigate adaptive governance models that evolve in response to system feedback, incorporating real-time ethical audits to mitigate inequities in primary care and beyond [16-18]. This includes longitudinal research on human-AI interaction dynamics, quantifying how fusion architectures impact clinical outcomes in diverse settings [19, 20].

Technological innovations offer promising avenues, such as integrating edge computing for real-time analytics in resource-limited environments, reducing latency in closed-loop systems [24, 27]. Research agendas should emphasize reproducibility and the development of standardized benchmarks for deep learning in EHR analysis and imaging to facilitate cross-system validation [23, 28]. In specialized domains such as mental health and neurology, future work could explore hybrid models that combine deep learning with causal inference to enable more robust prognostic analytics [6, 14].

Moreover, interdisciplinary collaborations should drive research into scalable deployment strategies, including simulation-based testing of infrastructures under varied clinical stressors [8, 10]. This could extend to policy-oriented studies that inform regulatory frameworks for AI analytics, ensuring alignment with global healthcare standards [5, 11]. Emerging directions also include the exploration of quantum-inspired deep learning for handling ultra-high-dimensional data, potentially revolutionizing analytics in oncology and gastroenterology [12].

To formalize one prospective direction, consider a governance-augmented feedback system: Let G denote governance protocols, integrated into the decision cycle as $D \to I(C) \to V \to F(G(D))$, where G applies ethical filters to updated data, promoting sustainable analytics evolution. This conceptual formula could guide research in designing resilient systems. Overall, these directions aim to propel deep learning toward mature, equitable integration, fostering proactive, patient-centric healthcare infrastructures [25, 26].

Conclusion

In conclusion, this systems-oriented review elucidates the profound impact of deep learning integration on clinical decision infrastructure, synthesizing a body of literature that highlights its role in advancing healthcare systems and analytics. From multimodal data fusion to closed-loop architectures, deep learning emerges as a cornerstone for intelligent, adaptive ecosystems that enhance clinical efficacy and equity. However, challenges in bias, scalability, and governance underscore the need for cautious, governed deployments.

Future research must bridge these gaps through innovative frameworks and interdisciplinary efforts, paving the way for resilient infrastructures that prioritize patient outcomes. Ultimately, the strategic embedding of deep learning promises a transformative era in healthcare, where analytics drive precise, ethical, and systemic decisions.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 19 Feb 2023 Revised: 29 Mar 2023 Accepted: 03 May 2023
Published online: 20 July 2023

References

Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: benefits, risks, and strategies for success. NPJ Digit Med. 2020;3(1):17. https://doi.org/10.1038/s41746-020-0221-y.

Huang SC, Pareek A, Seyyedi S, Banerjee I, Lungren MP. Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. NPJ Digit Med. 2020;3(1):136. https://doi.org/10.1038/s41746-020-00341-z.

Kline A, Wang H, Li Y, Dennis S, Hutch M, Xu Z, et al. Multimodal machine learning in precision health: A scoping review. NPJ Digit Med. 2022;5(1):171.
https://doi.org/10.1038/s41746-022-00712-8.

Aggarwal R, Sounderajah V, Martin G, Ting DS, Karthikesalingam A, King D, et al. Diagnostic accuracy of deep learning in medical imaging: a systematic review and meta-analysis. NPJ Digit Med. 2021;4(1):65.
https://doi.org/10.1038/s41746-021-00438-z.

Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33.
https://doi.org/10.1038/s41591-022-01772-9.

Su C, Xu Z, Pathak J, Wang F. Deep learning in mental health outcome research: a scoping review. Transl Psychiatry. 2020;10(1):116.
https://doi.org/10.1038/s41398-020-0780-3.

Liu X, Faes L, Kale AU, Chaturvedi S, Bourges D, Wagner SK, 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.

Ngiam KY, Khor IW. Big data and machine learning algorithms for health-care delivery. Lancet Oncol. 2019;20(5):e262-e273.
https://doi.org/10.1016/S1470-2045(19)30149-4.

Wang F, Casalino LP, Khullar D. Deep Learning in Medicine—Promise, Progress, and Challenges. JAMA Intern Med. 2019;179(3):293-4.
https://doi.org/10.1001/jamainternmed.2018.7117.

Jiang F, Jiang Y, Zhi H, Dong Y, Li H, Ma S, et al. Artificial intelligence in healthcare: past, present and future. Stroke Vasc Neurol. 2017;2(4):230.

Nagendran M, Chen Y, Lovejoy CA, Gordon AC, Komorowski M, Harvey H, et al. Artificial intelligence versus clinicians: systematic review of design, reporting standards, and claims of deep learning studies. BMJ. 2020;368:m689.
https://doi.org/10.1136/bmj.m689.

Gonçalves e Gonçalves W, de Paula dos Santos MH, França Lobato FM, Ribeiro-dos-Santos Â, Santana de Araújo G. Deep learning in gastric tissue diseases: a systematic review. BMJ Open Gastroenterol. 2020;7(1):e000371.

Ting DSW, Pasquale LR, Peng L, Campbell JP, Lee AY, Raman R, et al. Artificial intelligence and deep learning in ophthalmology. Br J Ophthalmol. 2019;103(2):167.
https://doi.org/10.1136/bjophthalmol-2018-313173.

Auger SD, Jacobs BM, Dobson R, Marshall CR, Noyce AJ. Big data, machine learning and artificial intelligence: a neurologist's guide. Pract Neurol. 2021;21(1).
https://doi.org/10.1136/practneurol-2020-002688.

Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. J Med Ethics. 2020;46(3):205.
https://doi.org/10.1136/medethics-2019-105586.

d'Elia A, Gabbay M, Rodgers S, Kierans C, Jones E, Durrani I, et al. Artificial intelligence and health inequities in primary care: a systematic scoping review and framework. Fam Med Community Health. 2022;10(Suppl 1):e001670.

Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. BMJ Qual Saf. 2019;28(3):231.
https://doi.org/10.1136/bmjqs-2018-008370.

Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368:l6927.
https://doi.org/10.1136/bmj.l6927.

Chen JH, Asch SM. Machine learning and prediction in medicine — beyond the peak of inflated expectations. N Engl J Med. 2017;376(26):2507-9.
https://doi.org/10.1056/NEJMp1702071.

Char DS, Shah NH, Magnus D. Implementing Machine Learning in Health Care — Addressing Ethical Challenges. N Engl J Med. 2018;378(11):981-3.
https://doi.org/10.1056/NEJMp1714229.

Babic B, Gerke S, Evgeniou T, Cohen IG. Beware explanations from AI in health care. Science. 2021;373(6552):284-6.
https://doi.org/10.1126/science.abg1834.

Ting DSW, Cheung CY, Lim G, Tan GSW, Quang ND, Gan A, et al. Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. JAMA. 2017;318(22):2211-23.
https://doi.org/10.1001/jama.2017.18152.

Wang S, Marinsek N, Ranganath R. Reproducibility in machine learning for health research: Still a ways to go. Sci Transl Med. 2021;13(586).
https://doi.org/10.1126/scitranslmed.abb1655.

Yasaka K, Abe O. Deep learning and artificial intelligence in radiology: Current applications and future directions. PLoS Med. 2018;15(11):e1002707.
https://doi.org/10.1371/journal.pmed.1002707.

Saria S, Butte A, Sheikh A. Better medicine through machine learning: What's real, and what's artificial? PLoS Med. 2018;15(12):e1002721.
https://doi.org/10.1371/journal.pmed.1002721.

Ashrafian H, Darzi A. Transforming health policy through machine learning. PLoS Med. 2018;15(11):e1002692.
https://doi.org/10.1371/journal.pmed.1002692.

Ravì D, Wong C, Deligianni F, Berthelot M, Andreu-Perez J, Lo B, et al. Deep learning for health informatics. IEEE J Biomed Health Inform. 2017;21(1):4-21.
https://doi.org/10.1109/JBHI.2016.2636665.

Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep EHR: a survey of recent advances in deep learning techniques for Electronic Health Record (EHR) analysis. IEEE J Biomed Health Inform. 2018;22(5):1589-604.
https://doi.org/10.1109/JBHI.2017.2767063.

Panayides AS, Amini A, Filipovic ND, Sharma H, Tsaftaris SA, Young A, et al. AI in medical imaging informatics: current challenges and future directions. IEEE J Biomed Health Inform. 2020;24(7):1837-57.
https://doi.org/10.1109/JBHI.2020.2991043.