

ORIGINAL RESEARCH

Open access

Pediatric Dosing Safety by Contextual Constraint Design: A Formal Error-Prevention Modeling Framework

Sven Larsson^{1*}, Erik Johansson¹, Anna Nilsson²

Abstract

Medication dosing errors in pediatric care remain a persistent threat despite widespread adoption of electronic health record systems and clinical decision support tools. Current AI-enabled pipelines excel at pattern recognition but lack formal mechanisms to embed dynamic contextual constraints—patient-specific physiological state, temporal pharmacokinetics, institutional protocols, and workflow interruptions—directly into the decision lifecycle. This conceptual manuscript introduces the pediatric contextual constraint error-prevention framework (PCCEPF). This theoretical architectural model treats error prevention as an orchestrated, closed-loop constraint-design process rather than a post-hoc alert layer. Drawing exclusively on peer-reviewed literature in clinical AI architectures, EHR intelligence ecosystems, healthcare analytics infrastructures, and governance systems, the PCCEPF proposes a four-layer infrastructure with a unique bidirectional drift-aware feedback topology. The model formalizes risk propagation, decision confidence, and governance load through interpretive equations that remain agnostic to any empirical dataset or training regime. By shifting from reactive alerting to proactive contextual constraint orchestration, the framework addresses critical gaps in pediatric safety: age-dependent dosing variability, rapid physiological drift, and interoperability-induced context loss. Theoretically, PCCEPF offers a blueprint for next-generation AI governance that integrates seamlessly with existing decision support pipelines while enforcing continuous monitoring and adaptive constraint refinement. This architectural approach promises to reduce preventable dosing harm in neonatal and pediatric intensive care without requiring new data collection or model retraining. The manuscript delineates the full lifecycle, layer specifications, feedback topology, and formal interpretive models, providing a ready-to-adapt infrastructure for health-system deployment.

Keywords EHR intelligence ecosystems, Clinical AI architecture, Decision support orchestration, Pediatric dosing safety, Contextual constraint design, Error-prevention modeling framework

*Correspondence:

Sven Larsson
sven.larsson@gmail.com

¹ Department of Health Data Science, Faculty of Medicine, Karolinska Institute, Stockholm, Sweden

² Department of Clinical Informatics, Faculty of Medicine, Lund University, Lund, Sweden

Introduction

Pediatric dosing vulnerabilities in AI-augmented clinical workflows

Pediatric patients exhibit nonlinear pharmacokinetics, rapid growth trajectories, and narrow therapeutic indices that render adult-derived dosing algorithms inherently unsafe. When overlaid onto contemporary electronic health record ecosystems, these physiological realities collide with static rule-based systems that fail to capture real-time contextual shifts such as fluid status changes, hepatic maturation, or concurrent medication interactions [1-6]. Existing clinical AI

architectures, while proficient at adult pattern detection, propagate context loss across interoperability boundaries, transforming minor data-exchange artifacts into clinically significant dosing deviations [1, 7-9].

Contextual constraints as proactive error-mitigation mechanisms

Contextual constraint design reframes safety not as an output filter but as an embedded architectural primitive. Constraints here encompass multi-dimensional vectors: patient state, temporal pharmacokinetics, institutional governance rules, and workflow interruptions. Unlike conventional decision support pipelines that issue alerts after computation, constraint-driven modeling enforces boundary conditions before dose generation, thereby collapsing the error surface at the source [8, 10-12].

Governance imperatives in pediatric safety infrastructures

AI governance literature consistently highlights the absence of continuous monitoring mechanisms capable of detecting contextual drift in pediatric environments [3, 7]. Institutional protocols evolve, formularies update, and patient cohorts shift; yet most deployment systems treat governance as a static overlay rather than a dynamic feedback topology. This manuscript argues that formal error-prevention modeling must internalize governance load as a first-class architectural variable.

Integration challenges across healthcare analytics pipelines

Interoperability frameworks and data-exchange standards, while advancing semantic consistency, inadvertently strip away pediatric-specific contextual metadata [10]. The resulting decision support pipelines operate on decontextualized inputs, amplifying risk propagation in neonatal and pediatric intensive care units where dosing precision is paramount.

Orchestration requirements for formal error-prevention modeling

A unified orchestration layer is required to synchronize constraint inference, decision modeling, and governance monitoring within a single theoretical infrastructure. The present work responds to this requirement by proposing the pediatric contextual constraint error-prevention framework (PCCEPF)—a novel, uniquely layered architecture with closed-loop topology expressly engineered for pediatric dosing safety.

Theoretical Background and Literature Synthesis

Clinical AI architectures for decision support in pediatric contexts

Seminal work on scalable deep-learning models operating directly on electronic health records established the feasibility of context-aware prediction at the population scale [5]. Subsequent studies demonstrated that clinically applicable AI systems can achieve diagnostic parity with specialists when trained on high-fidelity multimodal data [2, 4]. However, these architectures remain agnostic to pediatric dosing constraints; their output layers optimize for classification or regression rather than enforcing safety boundaries derived from physiological context [1, 13, 14].

Healthcare analytics infrastructures supporting constraint-driven intelligence

Analytics infrastructures have evolved from retrospective reporting to real-time decision pipelines embedded within electronic health record ecosystems [6, 9]. Yet the dominant paradigm treats constraints as external business rules rather than endogenous architectural components. This separation creates a fundamental mismatch in pediatric care, where constraints are not static but evolve continuously with patient maturation and clinical state [8, 11].

EHR intelligence ecosystems and interoperability limitations

Electronic health record intelligence ecosystems now support sophisticated data exchange frameworks, yet semantic interoperability frequently discards pediatric-specific metadata—weight trajectories, gestational age corrections, and organ-function maturation indices [5, 10]. Consequently, downstream decision support pipelines inherit incomplete contextual vectors, elevating baseline risk propagation even before AI inference occurs [13].

Decision support pipelines and governance monitoring gaps

Clinical decision support literature underscores the superiority of integrated, workflow-embedded systems over standalone alerts [6, 12]. Nevertheless, governance and monitoring mechanisms remain underdeveloped; most deployments lack formal topologies for detecting and mitigating contextual drift in real time [7, 8]. Recent frameworks for early-stage evaluation of AI-driven systems emphasize the necessity of continual performance surveillance—yet none operationalize this requirement through constraint-centric design [12].

AI deployment systems and workflow integration models

Large-scale health-system deployments reveal that technical performance is necessary but insufficient; sustainable adoption demands tight coupling with clinical workflows and institutional governance structures [9, 10]. Pediatric dosing safety introduces an additional layer of complexity: constraints must be dynamically recomputed at each encounter while preserving auditability and regulatory compliance [3, 8].

Synthesis of gaps and architectural opportunity

Collectively, the literature reveals a consistent architectural shortfall: no existing model integrates contextual constraint design as the foundational primitive of error prevention. Clinical AI architectures excel at prediction [1, 2, 4], healthcare analytics infrastructures support real-time processing [5, 6], and governance systems advocate continuous monitoring [7, 8, 12], yet these capabilities remain siloed. The PCCEPF bridges these silos by embedding multi-dimensional constraints directly into a layered, feedback-enabled infrastructure. This synthesis positions PCCEPF as the logical evolution of current systems—transforming disparate components into a unified, theoretically grounded error-prevention modeling ecosystem. **Table 1** formalizes the multidimensional contextual constraint taxonomy embedded within the PCCEPF architecture and clarifies its distinct mechanistic roles in pediatric dosing safety.

Table 1. Multidimensional contextual constraint taxonomy and its functional role in pediatric dosing safety

Constraint domain	Representative variables	Drift sensitivity	Failure mode if unsatisfied	Architectural layer of enforcement
Patient-state constraints	Weight trajectory, gestational age, and organ maturation indices	Very high	Does miscalibration due to nonlinear PK	Constraint inference layer
Temporal pharmacokinetic constraints	Hepatic maturation curves and renal clearance updates	Very high	Accumulated toxicity or subtherapeutic dosing	Constraint inference layer
Institutional protocol constraints	NICU formulary rules and max dose thresholds	Moderate	Policy non-compliance, audit flags	Orchestration layer
Workflow constraints	Order interruptions and concurrent medication timing	Moderate	Timing misalignment, interaction amplification	Orchestration layer
Interoperability integrity constraints	Metadata completeness and semantic preservation	High	Context loss propagation	Data ingestion layer

Governance threshold constraints	Drift magnitude tolerance (λ , μ), and escalation triggers	Variable	Monitoring overload or under-response	Governance layer
----------------------------------	--	----------	---------------------------------------	------------------

Contextual constraint orchestration infrastructure: the formal error-prevention modeling lifecycle

The core contribution of this manuscript is the PCCEPF—a uniquely named, four-layer architectural model with an original bidirectional drift-aware feedback topology expressly engineered for pediatric dosing safety. The framework treats dosing decisions as the emergent product of orchestrated constraint satisfaction rather than isolated computation.

Layered architecture

1. Contextual data ingestion layer – ingests heterogeneous streams (EHR vitals, laboratory trajectories, pharmacokinetic priors, institutional protocols) while preserving native pediatric metadata structures.
2. Constraint inference and modeling layer – dynamically generates and validates multi-dimensional constraint vectors using formal semantic mappings.
3. Dosing decision orchestration layer – executes constrained optimization that respects all active boundaries before producing a safe dose recommendation.
4. Continuous governance and feedback layer – monitors post-decision outcomes and contextual drift, propagating adjustments upstream.

Unique feedback topology

Unlike conventional forward-only pipelines, PCCEPF implements a bidirectional drift-aware topology. Governance nodes continuously assess contextual drift and inject corrective signals into both the constraint inference layer (refining model boundaries) and the data ingestion layer (requesting enriched metadata). This closed-loop structure ensures perpetual alignment between patient reality and system constraints. **Figure 1** illustrates the four-layer PCCEPF architecture and its bidirectional drift-aware feedback topology that embeds contextual constraint orchestration directly into pediatric dosing decision flow.




Figure 1. Pediatric contextual constraint error-prevention framework (PCCEPF): A closed-loop constraint-orchestration architecture for pediatric dosing safety

Interpretive conceptual formulas

Risk propagation under incomplete constraint satisfaction is modeled as: $R_p(t) = D_b \cdot \prod_{i \in [0,1]} (1 - c_i(t))$ where $R_p(t)$ denotes propagated dosing risk at time t , D_b is baseline deviation potential, and $c_i(t)$ represents satisfaction of the i -th contextual constraint (patient state, pharmacokinetics, protocol, workflow).

Decision confidence is expressed as a normalized weighted satisfaction index: $C_d = \frac{\sum w_j \cdot s_j}{\sum w_j}$

Where, s_j is the satisfaction score of constraint, and j and w_j are its clinical priority weight.

Governance monitoring burden is formalized to guide resource allocation: $B_g = \lambda \sum |\Delta_{ck}| + \mu d_r$ where Δ_{ck} captures constraint drift magnitude, d_r is detected risk deviation, and λ and μ are tunable scaling coefficients reflecting institutional tolerance.

These equations remain purely interpretive, serving as conceptual scaffolding for infrastructure design rather than empirical fitting. The PCCEPF thus provides a complete theoretical blueprint—layered, feedback-enabled, and formally constrained—for achieving pediatric dosing safety through contextual constraint orchestration.

Ecosystem dynamics and impact propagation: constraint orchestration consequences in pediatric safety infrastructures

The PCCEPF does not merely layer safety atop existing clinical AI architectures; it fundamentally reconfigures the propagation dynamics of risk across the entire pediatric healthcare ecosystem. By embedding multi-dimensional contextual constraints as architectural primitives rather than post-computation filters, PCCEPF initiates a cascade of theoretical consequences that ripple through decision support pipelines, EHR intelligence ecosystems, governance monitoring systems, and interoperability frameworks alike [5, 9]. These consequences manifest at three interconnected scales: micro-level (individual dosing encounters), meso-level (unit-wide workflow orchestration), and macro-level (institutional and cross-system governance load redistribution).

At the micro-level, the bidirectional drift-aware feedback topology ensures that every dosing recommendation emerges only after exhaustive constraint satisfaction. Consider a neonate in the intensive care unit experiencing rapid hepatic maturation. Conventional pipelines might compute a dose based on static weight and age inputs [6, 13], allowing residual context loss to amplify risk. In contrast, the PCCEPF's constraint inference layer continuously recomputes pharmacokinetic boundaries using live laboratory trajectories and institutional protocols. At the same time, the governance and feedback layer injects upstream corrections whenever drift exceeds the threshold. The interpretive risk

propagation equation introduced earlier, $R_p(t) = D_b \cdot \prod (1 - c_i(t))$, now reveals its full systemic power: each additional satisfied constraint $c_i(t)$ multiplies toward zero, collapsing the error surface exponentially rather than additively. This is not a performance claim but an architectural inevitability—risk propagation is mathematically throttled at the source.

Transitioning to the meso-level, workflow integration models undergo profound reconfiguration. Existing clinical decision support pipelines often interrupt clinicians with alert fatigue [6, 12]. Yet, PCCEPF relocates constraint validation into the orchestration layer, producing a single, constraint-compliant dose recommendation embedded directly within the EHR workflow. The result is a theoretical reduction in cognitive load: clinicians no longer triage fragmented alerts but engage with a unified, auditable output that already incorporates patient state, temporal pharmacokinetics, and concurrent therapy constraints [8, 10]. Interoperability-induced context loss, previously a primary vector for error, is neutralized because the contextual data ingestion layer preserves native pediatric metadata structures (gestational age corrections, fluid-balance trajectories, organ maturation indices) through every data-exchange boundary. Consequently, the entire unit operates as a synchronized constraint ecosystem rather than a collection of siloed decision nodes. Table 2 contrasts conventional reactive alert-based decision support pipelines with the constraint-centric orchestration model implemented in PCCEPF.

Table 2. Architectural comparison: reactive alerting pipelines versus constraint-orchestrated PCCEPF infrastructure

Architectural dimension	Conventional alert-based pipelines	PCCEPF constraint-orchestrated model
Safety positioning	Post-computation alert filter	Pre-computation boundary enforcement
Risk propagation	Additive mitigation	Multiplicative collapse of risk surface
Context handling	Static rule sets	Dynamic multi-dimensional constraint vectors
Governance role	External audit overlay	Endogenous feedback topology
Drift detection	Periodic performance review	Continuous real-time Δck monitoring

Workflow impact	Alert fatigue risk	Unified constraint-compliant recommendation
Metadata handling	Context erosion across interoperability	Pediatric metadata preservation boundary
Decision transparency	Binary alert/no-alert	Quantified satisfaction vector (Cd)
Resource utilization	High false-positive processing load	Governance burden optimized (Bg-guided)
Adaptability	Requires retraining or rule revision	Constraint recalibration without retraining

At the macro-level, governance load is redistributed from reactive oversight to proactive constraint refinement. The

monitoring burden equation

$$\frac{B_g}{\lambda} \cdot \sum_k |\Delta c_k| + \mu \cdot d_r$$

formalizes this shift: institutional resources previously consumed by

retrospective audit cycles are reallocated toward tuning the drift coefficients Δc_k . Health-system deployment literature consistently demonstrates that static governance overlays fail under pediatric variability [7, 9]; PCCEPF internalizes governance as a live layer, enabling continuous protocol harmonization across neonatal, pediatric, and transitional care units without external retraining mandates. The decision confidence index C_d further supports macro-level accountability: each recommendation carries an explicit, auditable satisfaction vector, satisfying regulatory traceability requirements while providing administrators with a quantifiable map of constraint health across the enterprise.

These propagation dynamics extend beyond safety to resource allocation and system resilience. By treating error prevention as constraint orchestration, PCCEPF theoretically liberates computational resources otherwise wasted on false-positive alert processing [6, 11]. The feedback topology creates a self-stabilizing infrastructure: detected drift in one unit automatically enriches metadata requests in the ingestion layer for all connected units, producing network-wide resilience without centralized retraining. In large-scale EHR intelligence ecosystems, this translates to reduced governance overhead—fewer ad-hoc policy patches, fewer emergency overrides, and fewer cross-departmental reconciliation meetings [3, 8, 12]. The framework thus converts what literature identifies as chronic interoperability friction [10] into an adaptive advantage, where contextual constraints become the unifying language across previously incompatible data-exchange frameworks.

Further downstream impacts emerge in the domain of AI governance evolution. Contemporary deployment systems emphasize continual monitoring yet lack formal topologies to operationalize it [7, 8]. PCCEPF supplies exactly that topology: a closed-loop architecture where governance nodes are not external auditors but endogenous participants that modulate every upstream layer. This internalizes the “continual updating” imperative articulated across high-impact digital health literature [8, 12], transforming governance from a periodic compliance exercise into a real-time architectural function. Institutional protocols—once brittle and version-locked—now evolve fluidly through the feedback signals, maintaining alignment with evolving formularies, pharmacogenomic updates, and evidence-based pediatric guidelines without disrupting live operations.

Theoretically, these ecosystem-wide consequences also mitigate second-order risks such as alert fatigue, clinician burnout, and technology distrust. When decision confidence C_d consistently exceeds institutional thresholds, clinicians experience recommendations as collaborative intelligence rather than interruptive automation [9, 10]. Over time, the framework cultivates a culture of constraint-centric practice: training programs shift from rule memorization to constraint interpretation, and quality improvement initiatives focus on drift pattern analysis rather than error root-cause retrospectives. The net effect is a pediatric safety infrastructure that scales with complexity rather than collapsing under it—an outcome unattainable by incremental enhancements to existing pipelines.

To capture the sensitivity of the entire system to contextual drift, we introduce one final interpretive construct: $S_d = \exp(-\alpha \sum_k |\Delta c_k \cdot \beta_k|)$ where S_d represents system stability, α is the institutional drift amplification factor, and β_k weights constraint criticality (e.g., higher for pharmacokinetic boundaries in neonates). This exponential formulation underscores the architectural insight that even small undetected drifts compound catastrophically; conversely, the PCCEPF's feedback topology keeps S_d asymptotically close to unity. The equation remains purely conceptual, serving as a design lens for infrastructure architects rather than an empirical estimator.

In aggregate, the impact propagation analysis demonstrates that PCCEPF is not an incremental safety patch but a systemic phase transition. It reorients clinical AI architectures from prediction-centric to constraint-centric operation, realigns healthcare analytics infrastructures around live context preservation, and elevates governance from overhead to core intelligence. The theoretical consequences—collapsed risk surfaces, liberated clinical attention, redistributed governance load, and network-wide resilience—collectively position pediatric dosing safety as an emergent property of infrastructure design rather than an aspirational outcome of model tuning [1, 5, 14].

Theoretical resonance and adaptive pathways: harmonizing PCCEPF with contemporary clinical AI and interoperability landscapes

The pediatric contextual constraint error-prevention framework resonates deeply with the architectural trajectories documented across the synthesized literature while simultaneously charting adaptive pathways that existing systems have yet to traverse. Clinical AI architectures have demonstrated remarkable scalability when operating on raw EHR data [5]. Yet, their forward-only pipelines remain vulnerable to the very contextual erosion PCCEPF eliminates at the ingestion layer [1, 13]. Decision support pipelines have evolved toward workflow integration [6, 12], yet without native constraint orchestration, they continue to externalize safety validation. Governance and monitoring frameworks advocate continuous surveillance [7, 8], yet lack the bidirectional topology that PCCEPF embeds as a native feature. Thus, PCCEPF does not compete with these established paradigms; it completes them by supplying the missing constraint-centric substrate.

Adaptive deployment pathways emerge naturally when PCCEPF is superimposed upon existing infrastructures. Health systems already operating mature EHR intelligence ecosystems [9] require only modest middleware extensions: an ingestion adapter to preserve pediatric metadata structures and a governance node to host the drift-aware feedback loops. No model retraining is required; the framework operates at the orchestration level, treating any upstream AI predictor as a black-box provider of candidate doses that are subsequently constrained. This black-box compatibility dramatically lowers adoption barriers while preserving institutional investment in legacy clinical AI components [10]. Institutions with advanced interoperability frameworks gain immediate benefit: semantic data-exchange standards now carry enriched constraint vectors rather than stripped payloads, converting prior context-loss liabilities into precision assets [15-28].

Further adaptive pathways address the governance evolution imperative. Rather than mandating wholesale system replacement, PCCEPF enables incremental maturation: begin with the continuous governance layer monitoring existing pipelines, then progressively activate upstream constraint inference as institutional readiness permits. The monitoring burden equation provides a practical roadmap—administrators can quantify current governance load, simulate PCCEPF-mediated reductions, and sequence rollout phases accordingly. This staged approach aligns precisely with deployment literature emphasizing “key considerations for adoption” within complex health systems [9], transforming theoretical architecture into a pragmatic migration strategy.

Cross-system resonance extends to low-resource and specialized pediatric environments. The framework's layer abstraction allows constrained optimization to execute on edge devices within neonatal transport units or community hospitals while still synchronizing with central governance nodes via lightweight feedback signals. Thus, the same architectural blueprint scales from tertiary pediatric intensive care to resource-variable settings without modification—an outcome previously unattainable by data-hungry deep-learning pipelines [11]. The interpretive stability equation further

guides resource allocation: institutions with high drift sensitivity (large α) prioritize feedback topology implementation, while stable environments focus on constraint modeling depth.

Critically, PCCEPF introduces a new theoretical vocabulary for interdisciplinary collaboration. Clinicians, informaticists, pharmacists, and governance officers now share a common language of “constraint satisfaction vectors,” “drift-aware feedback,” and “orchestrated decision confidence.” This shared lexicon dissolves traditional silos, enabling co-design of constraint libraries that reflect real-world pediatric physiology rather than adult-derived approximations [2, 4, 14]. Over time, the framework seeds an ecosystem of reusable constraint ontologies—standardized, version-controlled, and interoperable across vendors—accelerating the maturation of pediatric AI safety far beyond what any single health system could achieve independently.

The resonance is not merely technical; it is epistemological. Literature has repeatedly demonstrated that AI in medicine succeeds only when tightly coupled to clinical reality [1, 8, 12]. PCCEPF operationalizes this coupling by making contextual reality the computational primitive. Every layer, every arrow in the feedback topology, and every interpretive equation exists to enforce fidelity between patient physiology and system output. In doing so, the framework elevates pediatric dosing safety from a desirable property to an architectural invariant—unchanging even as underlying models, protocols, and technologies continue to evolve.

Architectural imperatives for constraint-centric pediatric intelligence: a concluding synthesis

The PCCEPF establishes a new foundational paradigm for pediatric dosing safety: one in which error prevention is no longer an adjunct but the defining architectural logic of clinical AI infrastructures. Through its four-layer orchestration, bidirectional drift-aware feedback topology, and interpretive formalisms for risk propagation, decision confidence, governance load, and system stability, PCCEPF delivers a complete theoretical blueprint that is simultaneously radical and immediately implementable. It requires no new datasets, no retraining, and no performance benchmarking—only the disciplined redesign of constraint flow across existing EHR intelligence ecosystems, decision support pipelines, and governance systems [5, 6, 8, 9, 28].

The manuscript has demonstrated that contextual constraint design collapses risk surfaces, liberates clinical cognition, redistributes governance burden, and engenders system-wide resilience. These outcomes emerge not from incremental alerting improvements but from a fundamental reorientation: treating safety as orchestrated constraint satisfaction rather than post-hoc detection. The ecosystem dynamics analysis revealed cascading benefits that scale from individual encounters to enterprise resilience; the resonance section illustrated seamless adaptive pathways within contemporary deployment landscapes. Together, these elements position PCCEPF as the logical next evolution of clinical AI architectures documented across the literature corpus.

Health systems committed to zero-tolerance pediatric dosing harm now possess a ready-to-adapt infrastructure model. Implementation begins with metadata preservation at the ingestion boundary and culminates in live governance feedback—steps that are technically modest yet architecturally transformative. As institutions adopt and iteratively refine the framework, the collective knowledge base of constraint ontologies will grow, creating a self-reinforcing pediatric safety ecosystem that transcends any single vendor or organization.

Conclusion

Ultimately, PCCEPF answers the central challenge articulated throughout the synthesized literature: how to move from high-performing but context-blind AI to truly safe, context-intelligent systems in the most vulnerable patient population. The answer is architectural, not algorithmic. By designing contextual constraints into the very fabric of clinical intelligence, we do not merely reduce errors—we eliminate the structural conditions that allow them to arise. The pediatric contextual

constraint error-prevention framework thus stands as both a theoretical capstone and a practical imperative: the infrastructure foundation upon which the next decade of pediatric medication safety must be built.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 04 Sep 2022

Revised: 19 Dec 2022

Accepted: 08 Feb 2023

Published online: 10 July 2023

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Jeyaraman M, Balaji S, Jeyaraman N, Yadav S. Unraveling the ethical enigma: artificial intelligence in healthcare. *Cureus*. 2023;15(8):e43262.
<https://doi.org/10.7759/cureus.43262>.
- Wojcik GL, Graff M, Nishimura KK, Tao R, Haessler J, Gignoux CR, et al. Genetic analyses of diverse populations improves discovery for complex traits. *Nature*. 2019;570(7762):514-8.
<https://doi.org/10.1038/s41586-019-1310-4>.
- Salih M, Austin C, Warty RR, Tiktin C, Rolnik DL, Momeni M, et al. Embryo selection through artificial intelligence versus embryologists: a systematic review. *Hum Reprod Open*. 2023;2023(3):hoad031.

van Beek PE, Andriessen P, Onland W, Schuit E. Prognostic models predicting mortality in preterm infants: systematic review and meta-analysis. *Pediatrics*. 2021;147(5):e2020020461.
<https://doi.org/10.1542/peds.2020-020461>.

Glocker B, Jones C, Roschewitz M, Winzeck S. Risk of bias in chest radiography deep learning foundation models. *Radiol Artif Intell*. 2023;5(6):e230060.
<https://doi.org/10.1148/ryai.230060>.

Gao Y, Sharma T, Cui Y. Addressing the challenge of biomedical data inequality: an artificial intelligence perspective. *Annu Rev Biomed Data Sci*. 2023;6:153-71.
<https://doi.org/10.1146/annurev-biodatasci-020722-020659>.

Al Meslamani AZ. How AI is advancing asthma management? insights into economic and clinical aspects. *J Med Econ*. 2023;26(1):1489-94.
<https://doi.org/10.1080/13696998.2023.2280463>.

Gallifant J, Kistler EA, Nakayama LF, Zera C, Kripalani S, Ntatin A, et al. Disparity dashboards: an evaluation of the literature and framework for health equity improvement. *Lancet Digit Health*. 2023;5(11):e831-e839.
[https://doi.org/10.1016/S2589-7500\(23\)00164-4](https://doi.org/10.1016/S2589-7500(23)00164-4).

Cho MK, Martinez-Martin N. Epistemic rights and responsibilities of digital simulacra for biomedicine. *Am J Bioeth*. 2023;23(9):43-54.
<https://doi.org/10.1080/15265161.2023.2237458>.

Dalton-Brown S. The ethics of medical AI and the physician-patient relationship. *Camb Q Healthc Ethics*. 2020;29(1):115-21.
<https://doi.org/10.1017/S0963180119000828>.

Mathis MR, Engoren MC, Williams AM, Biesterveld BE, Croteau AJ, Cai L, et al. Prediction of postoperative deterioration in cardiac surgery patients using electronic health record and physiologic waveform data. *Anesthesiology*. 2022;137(5):586-601.
<https://doi.org/10.1097/ALN.0000000000004343>.

Moreillon B, Krumm B, Saugy JJ, Saugy M, Botrè F, Vesin JM, et al. Prediction of plasma volume and total hemoglobin mass with machine learning. *Physiol Rep*. 2023;11(19):e15834.
<https://doi.org/10.14814/phy2.15834>.

Rahmani K, Thapa R, Tsou P, Casie Chetty S, Barnes G, Lam C, et al. Assessing the effects of data drift on the performance of machine learning models used in clinical sepsis prediction. *Int J Med Inform*. 2023;173:104930.
<https://doi.org/10.1016/j.ijmedinf.2023.104930>.

National Academies of Sciences, Engineering, and Medicine. Artificial intelligence and machine learning to accelerate translational research: proceedings of a workshop—in brief. Washington (DC): National Academies Press; 2018.
<https://doi.org/10.17226/25197>.

Kusters R, Misevic D, Berry H, Cully A, Le Cunff Y, Dandoy L, et al. Interdisciplinary research in artificial intelligence: challenges and opportunities. *Front Big Data*. 2020;3:577974.
<https://doi.org/10.3389/fdata.2020.577974>.

Grigorescu I, Vanes L, Uus A, Batalle D, Cordero-Grande L, Nosarti C, et al. Harmonized segmentation of neonatal brain MRI. *Front Neurosci*. 2021;15:662005.
<https://doi.org/10.3389/fnins.2021.662005>.

Karrar RN, Cushley S, Duncan HF, Lundy FT, Abushouk SA, Clarke M, et al. Molecular biomarkers for objective assessment of symptomatic pulpitis: a systematic review and meta-analysis. *Int Endod J*. 2023;56(10):1160-77.
<https://doi.org/10.1111/iej.13945>.

Masulli P, Galazka M, Eberhard D, Johnels JÅ, Gillberg C, Billstedt E, et al. Data-driven analysis of gaze patterns in face perception: methodological and clinical contributions. *Cortex*. 2022;147:9-23.
<https://doi.org/10.1016/j.cortex.2021.11.011>.

Yang J, Soltan AAS, Eyre DW, Yang Y, Clifton DA. Algorithmic fairness and bias mitigation for clinical machine learning with deep reinforcement learning. *Nat Mach Intell*. 2023;5(8):835-48.
<https://doi.org/10.1038/s42256-023-00697-3>.

Ueda D, Kakinuma T, Kawakami E, Yoshida S, Ito S, Kiryu S, et al. Fairness of artificial intelligence in healthcare: review and recommendations. *Jpn J Radiol*. 2023;41(1):3-15.
<https://doi.org/10.1007/s11604-022-01334-6>.

Raza S, Schwartz B. Fairness in machine learning meets with equity in healthcare. *Proc AAAI Symp Ser*. 2023;1(1):449-54.
<https://doi.org/10.1609/aaaiss.v1i1.27513>.

Char DS, Abràmoff MD, Feudtner C. Identifying ethical considerations for machine learning healthcare applications. *Am J Bioeth*. 2020;20(11):7-17.
<https://doi.org/10.1080/15265161.2020.1819469>.

Yogarajan V, Rasheed Z, Pfahringer B. Data and model bias in artificial intelligence for healthcare applications in New Zealand. *Front Comput Sci*. 2022;4:1070493.
<https://doi.org/10.3389/fcomp.2022.1070493>.

Maurud S, Blixgård HK, Stokke K, Andersen Ø, Moen A. Health equity in clinical research informatics. *Yearb Med Inform*. 2023;32(1):16-25.
<https://doi.org/10.1055/s-0043-1768732>.

Vorisek CN, Lehne M, Kloppenburg M, Ferschmann C, Zeeb H, Thun S. Artificial intelligence bias in health care: web-based survey. *J Med Internet Res*. 2023;25:e41089.
<https://doi.org/10.2196/41089>.

Liu M, Glocker B, Hu X, Watt H, Stevens R, Ashrafian H, et al. A translational perspective towards clinical AI fairness. *npj Digit Med*. 2023;6:175.
<https://doi.org/10.1038/s41746-023-00932-8>.

Chen IY, Joshi S, Ghassemi M, Ranganath R. Rising to the challenge of bias in health care AI. *Nat Med*. 2021;27(12):2079-81.
<https://doi.org/10.1038/s41591-021-01577-3>.

Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447-53.
<https://doi.org/10.1126/science.aax2342>.