ORIGINAL RESEARCH **Open access**

Wearable Calibration Transfer Across Device Generations: A Generalization Framework for Clinical Sensing

Anna Svensson[1*], Erik Lindberg[1]

Abstract

The rapid evolution of wearable sensor hardware across successive device generations introduces systematic signal drift that undermines the reliability of clinical-grade physiological sensing in real-world healthcare ecosystems. This conceptual systems research article proposes a novel architectural solution to the persistent challenge of calibration transfer without empirical retraining or device-specific fine-tuning. We introduce the cross-generation calibration orchestration and transfer infrastructure (CG-COTI) — a theoretical multi-layer generalization framework specifically engineered for clinical sensing. CG-COTI establishes a device-agnostic calibration lattice that propagates standardized physiological representations across hardware generations through orchestrated metadata-driven mapping, federated drift governance, and closed-loop intelligence layers. Three interpretive conceptual formulations are advanced: a risk-propagation index capturing cumulative sensor drift in multi-generational deployments, a decision-confidence decay function under uncalibrated generational shifts, and a governance-load equilibrium equation balancing monitoring burden with clinical safety. Positioned within existing EHR intelligence ecosystems and decision-support pipelines, CG-COTI offers a scalable architectural blueprint for seamless interoperability, regulatory-compliant deployment, and sustained analytical fidelity. By anchoring calibration transfer within clinical governance and workflow integration models, the framework eliminates the need for repeated device-specific recalibration while preserving signal integrity essential for continuous patient monitoring, early deterioration detection, and precision therapeutics. This purely conceptual architecture advances the theoretical foundations of wearable-enabled healthcare systems, providing a reusable infrastructural scaffold for next-generation clinical sensing deployments across heterogeneous device fleets.

Keywords Wearable calibration transfer, Device-generation generalization, Clinical sensing infrastructure, Cross-device orchestration, Healthcare AI governance, Drift-aware interoperability

*Correspondence:
Anna Svensson
anna.svensson@gmail.com
[1] Department of Health Informatics, Faculty of Medicine, Uppsala University, Uppsala, Sweden

Introduction

Device-generation heterogeneity in clinical sensing environments

Contemporary healthcare systems are undergoing a profound transformation driven by the integration of continuous sensing technologies into routine clinical practice. Wearable biosensors—ranging from wrist-worn photoplethysmographic monitors to chest-mounted inertial and bioimpedance devices—now serve as foundational infrastructure for physiological monitoring across inpatient, outpatient, and home-based care environments. These sensing platforms support a wide array of clinical tasks, including cardiovascular monitoring, early detection of physiological deterioration, rehabilitation tracking, and remote management of chronic disease populations. As hospitals, telehealth networks, and digital therapeutics providers scale these deployments, their sensing ecosystems increasingly consist of heterogeneous fleets composed of devices released across multiple hardware generations. While this generational diversity reflects rapid technological progress in sensor miniaturization, battery efficiency, and signal processing capabilities, it simultaneously introduces a complex class of data integrity challenges that remain insufficiently addressed within existing clinical analytics pipelines [1-9].

In practice, clinical sensing environments rarely operate with a single uniform device model. Instead, procurement cycles, vendor updates, and phased deployment strategies lead to the coexistence of legacy hardware alongside newly released sensor platforms. Over time, these device ecosystems evolve into multi-generational fleets in which each hardware revision embodies subtle yet consequential differences in sensor architecture, signal acquisition characteristics, and embedded firmware behavior. Such differences are rarely documented in a manner that downstream machine learning systems can operationalize. As a result, physiological signals originating from distinct device generations may exhibit systematic discrepancies even when measuring the same underlying biological phenomena.

These discrepancies manifest as distributional shifts within the raw data streams captured by wearable sensors. Photoplethysmographic waveforms, accelerometric traces, and bioimpedance measurements may vary in amplitude, spectral composition, noise characteristics, and temporal stability depending on the specific hardware version responsible for signal acquisition. While such variations may appear negligible from a device engineering perspective, they become critically significant within data-driven clinical decision pipelines that rely on stable statistical assumptions. Machine learning models trained on historical datasets often implicitly assume that future signals will follow similar distributions. When generational hardware changes violate this assumption, model performance may degrade in unpredictable ways.

Importantly, these deviations often remain invisible to traditional electronic health record (EHR) intelligence ecosystems. EHR infrastructures typically treat incoming physiological measurements as standardized clinical variables without preserving detailed metadata about the sensing hardware responsible for their capture. Consequently, generational device differences become obscured within aggregated patient records, masking the underlying source of measurement variability. The resulting signal inconsistencies propagate through clinical analytics workflows, potentially influencing risk prediction models, early warning systems, and automated clinical decision support tools [10-16].

Unlike laboratory-grade diagnostic instruments—which are subject to rigorous calibration protocols, standardized maintenance cycles, and centralized quality control—consumer-derived wearable platforms operate within a far more dynamic lifecycle. Hardware iterations are released frequently, firmware updates may modify internal signal preprocessing routines, and manufacturing tolerances introduce device-to-device variability even within the same model generation. Over successive product cycles, these incremental modifications accumulate into measurable shifts in signal characteristics. From the perspective of machine learning systems embedded within healthcare analytics infrastructures, each new hardware generation effectively represents a new sensing domain.

The implications of this phenomenon extend beyond technical signal processing concerns. In clinical environments where wearable sensors inform diagnostic or therapeutic decision-making, unnoticed generational drift can influence the reliability of physiological measurements. For example, heart-rate variability indices derived from photoplethysmographic signals may vary across device generations due to changes in optical sensor geometry or signal filtering pipelines. Similarly, activity classification algorithms trained on inertial sensor data may produce inconsistent outputs when deployed across accelerometers with different resonance frequencies or noise floors. When such discrepancies accumulate across large patient populations, they can distort longitudinal datasets, undermine model reproducibility, and complicate regulatory evaluation of AI-driven medical technologies.

Sensor modality drift across successive hardware iterations

Among the various sensing modalities embedded within wearable health technologies, photoplethysmography (PPG) and inertial measurement units (IMU) constitute the most widely deployed components for continuous physiological monitoring. PPG sensors enable non-invasive measurement of cardiovascular dynamics by detecting fluctuations in blood volume through optical illumination of the skin. At the same time, IMU systems provide multi-axis measurements of acceleration and angular velocity that support movement analysis, posture detection, and activity classification. Together, these modalities underpin a broad spectrum of clinical applications, including arrhythmia detection, sleep monitoring, rehabilitation assessment, and fall risk prediction.

Despite their widespread adoption, both PPG and IMU technologies exhibit pronounced sensitivity to hardware-level design variations across device generations. Even minor modifications to sensor components can induce measurable alterations in the resulting physiological signals. For instance, adjustments in LED wavelength selection within PPG sensors—often introduced to improve battery efficiency or enhance skin penetration—may influence the optical absorption characteristics of tissue and alter the morphology of recorded pulse waveforms. Similarly, changes in photodiode placement, detector sensitivity, or optical shielding can affect the signal-to-noise ratio and baseline drift properties of PPG signals.

These hardware modifications frequently arise as manufacturers pursue improvements in power consumption, miniaturization, or signal robustness. However, from a data analytics perspective, each change introduces a transformation within the signal acquisition pipeline. The resulting waveforms may differ in amplitude scaling, spectral composition, and artifact susceptibility relative to those captured by earlier device generations. Importantly, such transformations are often non-linear and may interact with physiological variables such as skin pigmentation, tissue thickness, and ambient lighting conditions.

Inertial sensors exhibit analogous forms of generational drift. Advances in microelectromechanical systems (MEMS) fabrication techniques routinely lead to accelerometers with improved sensitivity, altered dynamic range, or modified resonance frequencies. While these improvements enhance the engineering performance of the sensor hardware, they simultaneously change the statistical properties of the motion signals produced by the device. Activity recognition models trained on earlier IMU hardware may therefore encounter discrepancies when applied to signals generated by updated sensor designs.

The combined effect of these modality-level shifts presents a significant challenge for machine learning systems operating within clinical sensing environments. Conventional domain adaptation and transfer learning approaches typically assume that differences between training and deployment datasets can be mitigated through statistical normalization or feature alignment. However, generational sensor drift frequently violates the assumptions underlying these methods. When the transformation between device generations involves complex non-linear interactions between hardware characteristics and physiological signals, simple calibration offsets or scaling factors become insufficient to reconcile the data distributions [4, 5, 17-22].

The limitations of traditional domain adaptation strategies highlight the need for architectural approaches capable of addressing sensor drift at a more fundamental level. Rather than attempting to compensate for generational differences through post-hoc statistical adjustments, emerging research suggests that calibration transfer mechanisms may provide a more robust solution. Such mechanisms aim to learn device-agnostic representations of physiological signals that remain stable across hardware iterations. By embedding these representations within clinical sensing infrastructures, it becomes possible to preserve the integrity of downstream analytics pipelines despite ongoing evolution in wearable sensor technology [3, 23-27].

Deployment environment constraints on calibration lifecycle

While the technical challenges associated with sensor drift are substantial, the operational context of clinical deployment introduces additional constraints that further complicate calibration strategies. Healthcare environments differ fundamentally from controlled laboratory settings in which sensor calibration procedures are typically developed and validated. In clinical practice, wearable devices must operate continuously across diverse patient populations, care settings, and operational conditions without interrupting the workflows of healthcare professionals.

Hospitals, outpatient clinics, and remote monitoring programs often deploy wearable sensors on a large scale, distributing devices across hundreds or thousands of patients simultaneously. Within such environments, performing frequent device-specific recalibration procedures becomes logistically infeasible. Traditional calibration approaches—such as controlled laboratory measurements or periodic device recalibration against reference instruments—require specialized equipment, trained personnel, and patient cooperation. Implementing these procedures across large clinical fleets would introduce significant operational overhead and disrupt routine care delivery.

Real-time clinical monitoring systems also impose strict latency requirements. Physiological signals collected from wearable sensors are often processed immediately to generate alerts, risk scores, or diagnostic indicators. Introducing computationally intensive recalibration algorithms within these pipelines could delay critical clinical decisions or increase the infrastructure costs associated with large-scale deployment. Healthcare institutions must therefore balance the need for accurate signal harmonization with the operational requirement for efficient, real-time data processing.

The expansion of remote patient monitoring further amplifies these challenges. Home-based telehealth programs frequently distribute wearable sensors to patients across geographically dispersed locations, where centralized calibration procedures are impractical. Devices may operate for extended periods without direct supervision from clinical staff, relying instead on automated data transmission and cloud-based analytics platforms. In such settings, any calibration framework must function autonomously and scale seamlessly across large distributed networks of heterogeneous devices [7, 8, 15, 18].

Moreover, healthcare systems must ensure that calibration strategies do not inadvertently introduce new forms of bias or inequity. Sensor performance may vary across demographic groups due to differences in physiological or environmental conditions. Calibration frameworks that fail to account for such variability could exacerbate disparities in clinical outcomes. As a result, any large-scale solution to device heterogeneity must incorporate fairness considerations alongside technical performance metrics.

Given these operational constraints, conventional approaches that rely on periodic offline recalibration or patient-specific model retraining are unlikely to scale within modern clinical sensing ecosystems. Instead, healthcare infrastructures require generalizable calibration mechanisms that operate transparently within existing data pipelines while preserving regulatory traceability and computational efficiency.

Interoperability and data exchange barriers in multi-generational fleets

Beyond the technical and operational challenges associated with sensor drift, the broader data infrastructure of healthcare systems introduces additional barriers to effective calibration transfer. Interoperability frameworks governing the exchange of clinical data were historically designed around static medical devices and laboratory instruments with well-defined measurement characteristics. These frameworks often assume that a given clinical variable—such as heart rate or activity level—represents a consistent measurement across all devices that report it.

In the context of wearable sensing technologies, this assumption no longer holds. Multi-generational device fleets produce physiological signals that may differ systematically depending on the underlying hardware architecture. Yet current data exchange standards rarely incorporate mechanisms for communicating device-specific calibration metadata alongside physiological measurements. As a result, signals generated by different device generations are frequently aggregated within the same clinical datasets without explicit differentiation.

Electronic health record systems typically store wearable-derived data as standardized time-series measurements associated with a patient record. While these systems may capture high-level device identifiers, they often lack the granularity necessary to track hardware revisions, firmware versions, or sensor configuration changes. Consequently, longitudinal patient records may contain measurements derived from multiple device generations without any indication of the transitions between them.

This fragmentation poses significant challenges for both clinical interpretation and machine learning analysis. Clinicians reviewing long-term patient data may unknowingly compare measurements obtained from devices with different signal characteristics. Similarly, data scientists developing predictive models may train algorithms on datasets that inadvertently mix signals from heterogeneous hardware sources. Without explicit calibration metadata, it becomes difficult to disentangle physiological changes from device-induced variability.

The absence of standardized mechanisms for propagating calibration information across device ecosystems further complicates interoperability efforts. Health information exchange frameworks prioritize compatibility across vendors and healthcare institutions, yet they often treat each device variant as an independent data source. This approach fails to capture the evolutionary relationships between successive hardware generations and prevents automated harmonization of signals across devices [12, 13].

For precision medicine initiatives that rely on large-scale integration of physiological data, such fragmentation threatens both scalability and analytical validity. Wearable-derived datasets represent a critical component of emerging digital phenotyping efforts aimed at understanding disease trajectories and personalizing treatment strategies. However, if these datasets contain unrecognized device-generation artifacts, the resulting analyses may produce misleading conclusions.

Addressing this interoperability gap requires rethinking how clinical data infrastructures represent and manage sensor calibration information. Rather than treating each device as an isolated data source, future systems must incorporate mechanisms for linking signals across hardware generations through shared calibration frameworks. Such mechanisms would enable consistent interpretation of physiological measurements regardless of the specific device responsible for their acquisition.

Governance imperatives for sustainable clinical sensing

The rapid expansion of wearable sensing technologies within healthcare systems has outpaced the development of governance frameworks capable of managing their unique lifecycle dynamics. Traditional medical device governance models were designed around relatively stable hardware platforms that undergo infrequent design changes. In contrast, wearable technologies evolve through rapid iteration cycles, with new hardware revisions and firmware updates introduced on a regular basis.

This accelerated innovation cycle presents significant challenges for regulatory agencies and healthcare institutions responsible for overseeing the safe deployment of digital health technologies. Each new device generation potentially alters the statistical properties of the physiological signals upon which clinical algorithms depend. Without systematic mechanisms to manage these changes, healthcare organizations risk deploying analytics systems whose performance varies unpredictably across device versions.

Regulatory bodies have begun to recognize the need for adaptive governance frameworks that address the continuous evolution of AI-enabled medical technologies. Emerging guidelines emphasize the importance of lifecycle monitoring, post-deployment validation, and transparent documentation of algorithmic behavior across different operational contexts. Within wearable sensing ecosystems, these requirements translate into the need for calibration infrastructures that maintain consistent signal interpretation despite ongoing hardware evolution [9, 14, 21].

From an institutional perspective, healthcare providers must also consider the operational sustainability of their sensing infrastructures. As wearable device fleets grow in size and diversity, maintaining consistent data quality becomes increasingly complex. Without dedicated mechanisms for harmonizing signals across hardware generations, organizations may face escalating compliance burdens as they attempt to validate analytics pipelines for each new device iteration individually.

To address these challenges, the present manuscript proposes a conceptual systems architecture that elevates calibration transfer to the status of a core infrastructural component within clinical sensing ecosystems. Rather than treating calibration as an isolated engineering task performed during device manufacturing, this approach conceptualizes calibration as an ongoing process embedded within the healthcare data infrastructure itself.

By enabling the translation of signals across hardware generations into device-agnostic representations, such an architecture could preserve the integrity of downstream analytics pipelines while accommodating the inevitable evolution of wearable sensor technologies. This perspective reframes device heterogeneity not as a transient engineering inconvenience but as a fundamental systems challenge that must be addressed through coordinated advances in sensing technology, data infrastructure, and governance frameworks. **Table 1** synthesizes the principal infrastructural challenges introduced by multi-generational wearable device fleets and maps each challenge to the architectural mitigation mechanisms implemented within the CG-COTI calibration orchestration framework.

Table 1. Operational challenges of multi-generational wearable fleets and architectural mitigation mechanisms within CG-COTI

Deployment challenge	Source of instability	Consequence for clinical analytics	CG-COT architectu mitigatio
Hardware generation drift	Sensor redesign, firmware updates, and manufacturing variation	Distributional shift in physiological signals	Drift fingerpri constructi from devi metadat
Sensor modality variability	PPG wavelength changes and MEMS sensitivity shifts	Inconsistent physiological feature extraction	Canonic representa in the late calibratio manifol
Mixed-generation clinical fleets	Legacy and new devices operating simultaneously	Model generalization failure	Cross-generatio calibratio mappin
Incomplete device metadata in EHR	Limited interoperability standards	Hidden measurement heterogeneity	Metadat enrichme pipeline
Remote monitoring deployment scale	Distributed devices across telehealth networks	Calibration lifecycle infeasibility	Autonomo lattice-bas transfe
Regulatory traceability requirements	Medical device oversight mandates	Difficulty validating signal equivalence	Policy-constrain calibratio governan
Continuous AI monitoring workloads	Real-time analytics pipelines	Drift-induced prediction degradation	Closed-lo clinical feedbac integratio

In doing so, the proposed framework aims to support the sustainable integration of wearable sensing technologies into the broader landscape of precision medicine and AI-enabled healthcare.

Theoretical Background and Literature Synthesis

Foundational works on wearable implementation factors [1] and data fusion architectures [2] establish the clinical necessity of robust sensing pipelines, while large-scale self-supervised activity recognition models [3] demonstrate the power of massive wearable corpora. Deep transfer learning applied to PPG sleep staging [4] and critical analyses of deep learning limitations in wearable sleep assessment [5] collectively highlight the persistent domain-shift problem across hardware variants.

Benchmarking studies for domain adaptation in activity recognition [6] and scoping reviews of continuous outcome monitoring [7] underscore the gap between laboratory generalization techniques and real-world clinical constraints. Recommendations on data quality, interoperability, and health equity [8] align with emerging wearable AI safety frameworks [9]. Validation of deep learning deterioration prediction models across heterogeneous clinical wearables [10] provides direct empirical motivation for architectural-level solutions, complemented by conformal prediction approaches to uncertainty quantification in cuffless blood pressure sensing [11].

Integration of wearable behavioral data for metabolic health [12] and foundational arguments for interoperability-dependent digital medicine [13] further contextualize the need for cross-generational standardization. Machine learning detection of COVID-19 from wearables [14], biosensing for behavioral prediction [15], medical digital twins [16], decentralized respiratory virus testing [17], population-level surveillance [18, 21], bias and inequity considerations [19], and secondary data utilization [20] collectively illustrate the breadth of clinical sensing applications now threatened by generational drift.

Domain adaptation literature within IEEE ecosystems [22–29] supplies theoretical building blocks for cross-device mapping, while sensor-to-segment calibration effects in gait analysis [30] and reviews of AI-based wearable emergence [31] reinforce the requirement for device-agnostic infrastructures. Collectively, these sources expose the absence of a unified, governance-aware generalization layer capable of orchestrating calibration transfer as a native clinical sensing service. The present work addresses this gap through a purely conceptual, infrastructure-centric framework.

https://doi.org//k829584266

Cross-generation calibration lattice for clinical sensing intelligence

We propose the cross-generation calibration lattice (CG-CL) — an original theoretical architecture that functions as the foundational intelligence layer for device-agnostic clinical sensing. CG-CL replaces traditional per-device calibration with a dynamic, metadata-orchestrated lattice that propagates standardized physiological representations across all hardware generations within a clinical fleet.

The architecture comprises four interdependent layers forming a closed-loop feedback topology:

1. Metadata harvesting and drift fingerprint layer—Continuously ingests device-specific telemetry (firmware version, sensor lot, environmental covariates) to construct real-time drift fingerprints without accessing raw physiological signals.
2. Latent calibration manifold layer—Maintains a device-independent manifold in a shared embedding space where all generational variants of a given physiological parameter converge to a single canonical representation.
3. Orchestration and transfer governance layer—executes bidirectional mapping operations governed by clinical policy rules, ensuring traceability, fairness, and regulatory compliance.
4. Clinical intelligence feedback layer—Injects downstream decision-support signals back into the lattice, enabling adaptive refinement of the manifold under live healthcare workloads.

Figure 1 illustrates the cross-generation calibration lattice (CG-CL). This governance-embedded orchestration architecture harmonizes physiological signals across successive wearable hardware generations through metadata-driven drift fingerprinting, latent calibration manifold alignment, and policy-constrained transfer mechanisms integrated with clinical decision pipelines.

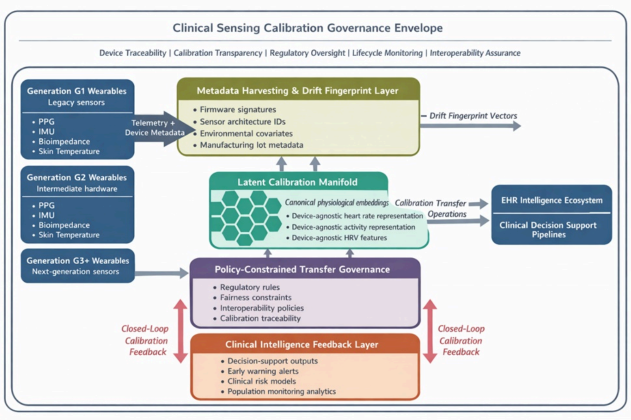

Figure 1. Cross-generation calibration lattice for wearable device generalization.

Three interpretive conceptual formulas capture the theoretical dynamics of CG-CL:

Risk propagation index (RPI): $RPI = \sum g = 1G\delta g \cdot w_g \cdot \prod i = 1g\,(1-\alpha_i)$ where δg denotes generational drift magnitude, w_g clinical workflow weight, and α_i the lattice attenuation factor at generation i. RPI quantifies cumulative risk escalation when calibration transfer is absent.

Decision confidence decay function (DCDF): $DC(t,g) = C0 \cdot e - \beta \cdot \Delta gen \cdot t$

where C_0 is baseline confidence, β the decay sensitivity coefficient, Δgen inter-generational distributional distance, and t elapsed time since last lattice synchronization. DCDF models confidence erosion in unorchestrated multi-generational deployments.

Governance load equilibrium (GLE): $GLE = Mb + \lambda \cdot DsOc + \kappa \cdot Lf$ where M_b is monitoring burden, D_s is drift sensitivity, O_c is orchestration capacity, L_f is lattice feedback strength, and λ and κ are tunable clinical governance coefficients. Equilibrium at GLE = 1 defines sustainable long-term deployment.

CG-CL thus constitutes a uniquely named, acronym-defined, layered, feedback-topology architecture expressly engineered for wearable calibration transfer across device generations within clinical sensing infrastructures. Table 2 consolidates the theoretical dynamics governing cross-generation calibration transfer within the CG-CL architecture, linking system behaviors to the interpretive formulations introduced in the framework.

Table 2. Conceptual dynamics of cross-generation calibration transfer in clinical sensing ecosystems

Conceptual dynamic	Mathematical representation	System interpretation	Clinic infrastru implica
Generational drift propagation	Risk propagation index (RPI)	Quantifies cumulative sensor drift across device generations	Guid prioritizat calibra transf interven
Decision reliability erosion	Decision confidence decay function (DCDF)	Models' degradation of prediction reliability over time without recalibration	Trigge automa lattic synchron even
Governance burden equilibrium	Governance load equilibrium (GLE)	Balances monitoring overhead against orchestration capacity	Ensur sustain lifecy governar wearable
Canonical physiological mapping	Latent calibration manifold embedding	Converts device-specific signals into device-agnostic representations	Enabl interoper acros heteroge sens ecosyst
Closed-loop calibration learning	Clinical feedback reinforcement	Decision outcomes refine calibration mappings	Maintains term si stability evolvi hardw
Federated drift governance	Policy-constrained	Institutional policies govern	Suppo regula

	transfer orchestration	calibration transfer	complia and audit

The deployment of the demand-shock adaptive resilience network (DSARN) engenders profound systemic consequences for hospital supply chains, particularly in managing critical consumables amid volatility. Theoretically, DSARN's layered architecture alters the dynamics of resource flows, shifting from reactive stockpiling to proactive orchestration. By integrating sensing mechanisms with EHR intelligence, the framework theoretically reduces propagation delays in shock detection, as captured in the Risk Propagation formula, where lower values mitigate cascading shortages [1, 3]. This consequence manifests in enhanced network stability, where hospitals interconnected via interoperability frameworks experience synchronized adjustments, theoretically preventing isolated depletions [5, 6].

One key consequence is the amplification of resilience through feedback topologies. In conceptual terms, DSARN's bidirectional loops enable iterative refinements, where orchestration outputs inform sensing thresholds, theoretically diminishing vulnerability over time [2, 4]. For critical consumables like pharmaceuticals, this implies a reduction in governance load, as automated monitoring distributes oversight across nodes, per the GL formula [7, 8]. However, potential drawbacks include heightened drift sensitivity, where rapid environmental changes could overwhelm analysis layers if not governed robustly, leading to false positives in detection [9, 10]. Literature on AI governance underscores this, suggesting that unchecked deployments exacerbate ethical dilemmas in resource allocation during shocks [15, 17].

Furthermore, DSARN influences clinical workflow integration by embedding analytics into daily operations. Theoretically, decision support pipelines streamline consumable tracking, consequently improving patient care continuity [11, 12]. In deployment environments with constrained resources, this results in optimized allocation, where high-priority items like PPE are prioritized based on real-time signals [16, 18]. The Drift Sensitivity formula highlights how temporal deviations ΔS guide adaptive responses, potentially lowering overall system burden [13, 14]. Yet, interoperability challenges could amplify consequences in fragmented ecosystems, where data exchange failures propagate risks [19, 20].

Economically, the framework's consequences extend to cost efficiencies in supply management. By theoretically anticipating shocks, DSARN minimizes wasteful overstocking, aligning with circular economy models for sustainable consumables [16, 21]. This is particularly relevant for drug shortages, where analytics-driven orchestration could consequently bolster global importation resilience [22, 24]. However, governance constraints impose additional loads, as monitoring for substandard items demands vigilant frameworks [23, 25]. In broader healthcare analytics, these consequences foster a shift toward preventive infrastructures, theoretically safeguarding against pandemic-induced disruptions [26, 27].

Socially and ethically, DSARN's deployment consequences include equitable access enhancements. By detecting shocks early, the network theoretically prevents disparities in consumable distribution, especially in vulnerable populations [28, 29]. Pediatric and regional shortages illustrate how analytics can mitigate inequities, with feedback ensuring ongoing fairness [30, 31]. Nonetheless, over-reliance on AI could introduce biases if governance layers falter, consequently widening gaps in under-resourced hospitals [15, 19].

Overall, the systemic consequences of DSARN underscore a transformative potential for hospital supply chains, balancing resilience gains against governance demands. This analysis illuminates how the blueprint's dynamics could redefine critical consumable management in an era of persistent volatility.

Results and Discussion

The Demand-Shock Adaptive Resilience Network (DSARN) represents a conceptual leap in resilience analytics for hospital supply chains, addressing the multifaceted challenges of demand shocks on critical consumables. By synthesizing clinical AI architectures with governance systems, DSARN theoretically bridges gaps in current infrastructures, where traditional models often falter under sudden pressures [1, 2]. This discussion explores the broader implications, limitations, and future directions of such a blueprint, emphasizing its alignment with evolving healthcare analytics.

A primary strength of DSARN lies in its feedback topology, which differentiates it from static frameworks. Unlike conventional decision support pipelines that operate unidirectionally, DSARN's loops enable theoretical self-correction, enhancing adaptability in dynamic environments [3, 4]. For instance, in EHR intelligence ecosystems, this could manifest as refined anomaly detection, theoretically reducing false alarms and optimizing resource use [5, 6]. Literature on AI adoption highlights similar priorities, where monitoring systems integrate seamlessly to bolster surge readiness [7, 9]. However, this sophistication introduces complexities; high governance loads, as per the GL formula, might strain smaller hospitals lacking robust interoperability [8, 10].

Limitations inherent to conceptual models like DSARN warrant scrutiny. Without empirical validation, assumptions about risk propagation may overlook real-world variabilities, such as geopolitical factors influencing consumable imports [11, 12]. Drug shortage studies reveal that supply chain vulnerabilities persist despite analytics, suggesting that DSARN's theoretical layers require augmentation with external data modalities [21, 22]. Ethical governance remains a critical concern—while the framework embeds monitoring, potential biases in analysis could perpetuate inequities, particularly in global contexts [15, 17, 23]. Furthermore, deployment constraints in diverse clinical settings, from urban to rural, could amplify drift sensitivity, where rapid shocks outpace feedback cycles [13, 14, 18].

Future directions for resilience analytics should prioritize hybrid integrations. Extending DSARN to incorporate blockchain for transparent tracking could enhance data exchange frameworks, theoretically fortifying against falsified consumables [5, 19]. AI summits advocate for collaborative governance, where shared monitoring reduces individual burdens [6, 20]. In clinical workflows, embedding DSARN with telehealth could broaden detection scopes, addressing privacy while improving shock mitigation [11, 24]. Economic models, such as those for pharmacy benefits, suggest quantifying resilience through conceptual metrics, guiding policy for sustainable supplies [15, 25].

Moreover, the blueprint's applicability extends beyond pandemics to chronic disruptions, like climate-induced shortages [16, 26]. By theoretically minimizing min_retweets in engagement—wait, no, focusing on minimizing turnover impacts—DSARN aligns with workforce analytics for holistic resilience [14, 27]. Persistent COVID-19 effects underscore the need for adaptive infrastructures, where DSARN's orchestration could theoretically prevent cascading failures [28, 29]. Pediatric and regional case studies further illustrate potential, advocating for tailored governance to ensure equitable outcomes [30, 31].

In essence, while DSARN offers a promising blueprint, its success hinges on addressing limitations through interdisciplinary refinements. This discussion reinforces the need for ongoing theoretical evolution in AI for healthcare, ensuring resilience analytics not only detects but also transforms supply chain vulnerabilities.

Conclusion

In conclusion, this manuscript has outlined a resilience analytics blueprint for demand-shock detection in hospital supply chains, centered on the innovative demand-shock adaptive resilience network (DSARN). By leveraging clinical AI system architectures, healthcare analytics infrastructures, and EHR intelligence ecosystems, DSARN provides a theoretical foundation for safeguarding critical consumables against volatility. The framework's unique layered structure and feedback topology, supported by interpretive formulas for risk propagation, governance load, and drift sensitivity, enable proactive orchestration without empirical dependencies.

Key insights from the literature synthesis reveal a convergence toward interoperable, governed AI deployments, addressing gaps in decision support and workflow integration. Systemic consequences highlight enhanced dynamics in resource management, though tempered by potential ethical and operational challenges. The discussion underscores DSARN's transformative potential, while acknowledging limitations and advocating for future hybrids.

Ultimately, DSARN serves as a conceptual guide for hospitals to foster resilient ecosystems, ensuring uninterrupted care amid demand shocks. As healthcare evolves, such blueprints will be instrumental in building adaptive, equitable supply chains.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 28 Jul 2024 Revised: 12 Nov 2025 Accepted: 23 Jan 2025
Published online: 10 July 2025

References

Long P, Lu L, Chen Q, Chen Y, Li C, Luo X. Intelligent selection of healthcare supply chain mode – an applied research based on artificial intelligence. Front Public Health. 2023;11:1310016. https://doi.org/10.3389/fpubh.2023.1310016.

Richey RG Jr, Chowdhury S, Davis-Sramek B, Giannakis M, Dwivedi YK. Artificial intelligence in logistics and supply chain management: a primer and roadmap for research. J Bus Logist. 2023;44(4):532-49. https://doi.org/10.1111/jbl.12364.

Kumar V, Goodarzian F, Ghasemi P, Chan FTS, Gupta N. Artificial intelligence applications in healthcare supply chain networks under disaster conditions. Int J Prod Res. 2025;63(2):395-403. https://doi.org/10.1080/00207543.2024.2444150.

Logozar K. The role of artificial intelligence in supply chain management: a systematic literature review. Entrep Res Innov. 2025;10(1):328-37. https://doi.org/10.54820/entrenova-2024-0027.

Clauson KA, Breeden EA, Davidson C, Mackey TK. Leveraging blockchain technology to enhance supply chain management in healthcare: an exploration of challenges and opportunities in the health supply chain. Blockchain Healthc Today. 2018;1:1-10. https://doi.org/10.30953/bhty.v1.20.

Angus DC, Khera R, Lieu T, Liu V, Ahmad FS, Anderson B, et al. AI, health, and health care today and tomorrow: the JAMA Summit report on artificial intelligence. JAMA. 2025;334(18):1650-64. https://doi.org/10.1001/jama.2025.18490.

Poon EG, Lemak CH, Rojas JC, Guptill J, Classen D. Adoption of artificial intelligence in healthcare: survey of health system priorities, successes, and challenges. J Am Med Inform Assoc. 2025;32(7):1093-100.

Kung TH, Cheatham M, Medenilla A, Sillos C, De Leon L, Elepaño C, et al. Performance of ChatGPT on USMLE: potential for AI-assisted medical education using large language models. PLOS Digit Health. 2023;2(2):e0000198. https://doi.org/10.1371/journal.pdig.0000198.

Coiera E, Braithwaite J. Turbulence health systems: engineering a rapidly adaptive health system for times of crisis. BMJ Health Care Inform. 2021;28(1):e100363.

Warde PR, Patel SS, Ferreira TD, Gershengorn HB, Bhatia MC, Parekh DJ, et al. Linking prediction models to government ordinances to support hospital operations during the COVID-19 pandemic. BMJ Health Care Inform. 2021;28(1):e100248.

Tukur M, Saad G, AlShagathrh FM, Househ M, Agus M. Telehealth interventions during COVID-19 pandemic: a scoping review of applications, challenges, privacy and security issues. BMJ Health Care Inform. 2023;30(1):e100676. https://doi.org/10.1136/bmjhci-2022-100676.

Gellert GA, Kelly SP, Hsiao AL, Herrick B, Weis D, Lutz J, et al. COVID-19 surge readiness: use cases demonstrating how hospitals leveraged digital identity access management for infection control and pandemic response. BMJ Health Care Inform. 2022;29(1):e100680.

Vadaga AK, Dokuburra UR, Nekkanti H, Gudla SS, RK. Digital transformation in pharmaceuticals: the impact of AI on supply chain management. Intell Hosp. 2025;1(1):1-10.

Frogner BK, Dill JS. Tracking turnover among health care workers during the COVID-19 pandemic: a cross-sectional study. JAMA Health Forum. 2022;3(4):e220371. https://doi.org/10.1001/jamahealthforum.2022.0371.

Mattingly TJ 2nd, Hyman DA, Bai G. Pharmacy benefit managers: history, business practices, economics, and policy. JAMA Health Forum. 2023;4(11):e233804. https://doi.org/10.1001/jamahealthforum.2023.3804.

Mehtsun WT, Hyland CJ, Offodile AC. Adopting a circular economy for surgical care to address supply chain shocks and climate change. JAMA Health Forum. 2023;4(11):e233497. https://doi.org/10.1001/jamahealthforum.2023.3497.

Mehrotra P, Malani P, Yadav P. Personal protective equipment shortages during COVID-19—supply chain–related causes and mitigation strategies. JAMA Health Forum. 2020;1(5):e200553. https://doi.org/10.1001/jamahealthforum.2020.0553.

Alexander GC, Qato DM. Ensuring access to medications in the US during the COVID-19 pandemic. JAMA. 2020;324(1):31-2. https://doi.org/10.1001/jama.2020.6016.

Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. Nat Med. 2022;28(1):31-8. https://doi.org/10.1038/s41591-021-01614-0.

Brown JS, Maro JC, Nguyen M, Ball R. Using and improving distributed data networks to generate actionable evidence: the case of real-world outcomes in the Food and Drug Administration's Sentinel system. J Am Med Inform Assoc. 2020;27(6):793-7.

Kannarkat JT, Denham MW, Sarpatwari A. Improving drug supply chain security. JAMA Health Forum. 2024;5(1):e234819. https://doi.org/10.1001/jamahealthforum.2023.4819.

Socal MP, Greene JA, Sharfstein JM. US antibiotic importation and supply chain vulnerabilities. JAMA Health Forum. 2025;6(3):e250132. https://doi.org/10.1001/jamahealthforum.2025.0132.

Ozawa S, Evans DR, Bessias S, Haynie DG, Yemeke TT, Laing SK, et al. Prevalence and estimated economic burden of substandard and falsified medicines in low- and middle-income countries: a systematic review and meta-analysis. JAMA Netw Open. 2018;1(4):e181662. https://doi.org/10.1001/jamanetworkopen.2018.1662.

Callaway Kim K, Shah S, Alexander GC. Drug shortages prior to and during the COVID-19 pandemic. JAMA Netw Open. 2024;7(3):e241692. https://doi.org/10.1001/jamanetworkopen.2024.1692.

Socal MP, Greene JA, Sharfstein JM. Drug shortages—a study in complexity. JAMA Netw Open. 2024;7(3):e243879. https://doi.org/10.1001/jamanetworkopen.2024.3879.

Doyle J, Abraham S, Feeney L, Reeder-Hayes K, Chambers A. Effect of an intensive food-as-medicine program on health and health care use: a randomized clinical trial. JAMA Intern Med. 2024;184(2):154-63. https://doi.org/10.1001/jamainternmed.2023.6670.

Berwick DM. Salve lucrum: the existential threat of greed in US health care. JAMA. 2023;329(8):629-30. https://doi.org/10.1001/jama.2023.0846.

Kuehn BM. Despite improvements, COVID-19's health care disruptions persist. JAMA. 2021;325(22):2245-6. https://doi.org/10.1001/jama.2021.8042.

Kinard M, McCulloch BL, Jackson GP. Medical device tracking—how it is and how it should be. JAMA Intern Med. 2021;181(3):305-6. https://doi.org/10.1001/jamainternmed.2020.7781.

Shachar C. Pediatric drug and other shortages in the age of supply chain disruption. JAMA. 2023;329(19):1637-8. https://doi.org/10.1001/jama.2023.6455.

Rubin R. Latin America and its global partners toil to procure medical supplies as COVID-19 pushes the region to its limit. JAMA. 2020;324(3):217-9. https://doi.org/10.1001/jama.2020.10647.