ORIGINAL RESEARCH Open access

Alarm Fatigue Mitigation Under Safety Constraints: A Context-Aware Suppression Design Framework

Ahmed Youssef[1*], Khaled Hassan[1]

Abstract

Alarm fatigue in healthcare settings poses significant risks to patient safety, arising from excessive, non-actionable alerts that desensitize clinicians. This conceptual manuscript introduces a novel framework for mitigating alarm fatigue through context-aware suppression mechanisms, while rigorously adhering to safety constraints. Drawing on theoretical principles from systems engineering, human factors, and artificial intelligence, we propose the safety-integrated context-aware suppression topology (SICAST), a multi-layered architecture designed to dynamically filter alarms based on real-time contextual data such as patient physiology, environmental factors, and clinician workload. The framework incorporates feedback loops for continuous adaptation, ensuring suppression decisions prioritize risk minimization without compromising vigilance. Key components include a context aggregation layer, a suppression decision engine governed by safety thresholds, and an audit trail for governance. Interpretive formulas model risk propagation under suppression and decision confidence amid constraints. By synthesizing recent literature, we highlight how SICAST addresses gaps in existing approaches, such as static thresholding and a lack of contextual integration. This work advances conceptual designs for AI-driven healthcare systems, emphasizing infrastructural resilience and ethical deployment. Implications for system orchestration in critical care underscore the need for balanced alarm management to enhance patient outcomes and reduce clinician burden.

Keywords Alarm fatigue, Context-aware suppression, Safety constraints, Healthcare alarm systems, AI mitigation framework, Risk propagation modeling

*Correspondence:
Ahmed Youssef
ahmed.youssef@gmail.com
[1] Department of Health Informatics, Faculty of Medicine, University of Khartoum, Khartoum, Sudan

Introduction

In contemporary healthcare environments, the proliferation of monitoring devices has led to an exponential increase in auditory and visual alerts, culminating in a phenomenon known as alarm fatigue. This condition, characterized by diminished responsiveness to alarms due to overload, undermines the foundational goals of patient monitoring systems. The manuscript explores mitigation strategies through a context-aware suppression design, ensuring that safety constraints are not violated amid efforts to reduce non-essential notifications. By conceptualizing alarm suppression as a dynamic, intelligent process, we aim to reconcile the tension between alert abundance and clinical efficacy.

Alarm overload dynamics in intensive care settings

Critical care units exemplify the alarm fatigue challenge, where multiple devices generate thousands of alerts daily, many of which are false positives or clinically insignificant [1, 2]. This overload stems from default settings that prioritize sensitivity over specificity, leading to sensory bombardment that erodes clinician attention. Contextual elements, such as shift timing and patient acuity, exacerbate the issue, as alarms fail to adapt to varying environmental demands. Literature indicates that up to 94% of alarms in intensive care may be non-actionable, fostering a cycle of desensitization that correlates with delayed responses to genuine emergencies [3, 4]. Safety constraints demand that any mitigation approach preserve the detection of critical events, yet traditional systems lack the intelligence to discern relevance based on holistic context.

Safety ramifications of unsuppressed alarm proliferation

Unmanaged alarm fatigue directly imperils patient safety, manifesting in adverse events such as missed deteriorations or medication errors [5, 6]. When clinicians are exposed to excessive alarms—many of which may be clinically insignificant—their responsiveness to critical alerts may diminish over time, increasing the probability that meaningful warnings are delayed or ignored. Under safety constraints, suppression designs must incorporate robust fail-safe mechanisms to prevent alert omission, carefully balancing alarm reduction with reliability and clinical visibility. Systems that suppress alerts without rigorous safeguards risk introducing new safety hazards, particularly in high-acuity environments where rapid detection of physiological deterioration is essential.

Theoretical models suggest that fatigue-induced errors arise from cognitive overload, where clinicians' decision-making capacity is compromised by incessant interruptions and fragmented attention [7, 8]. Continuous alarm exposure may disrupt situational awareness, forcing healthcare professionals to divide attention across competing tasks. In surgical or emergency contexts, this phenomenon can amplify risks, as alarms compete with procedural demands, time-critical decision making, and coordination among care teams. The manuscript posits that context-aware mechanisms can mitigate these ramifications by integrating real-time clinical data streams, ensuring suppression only occurs when safety thresholds are clearly satisfied and when the probability of clinically meaningful deterioration remains low.

Contextual modalities shaping alarm relevance

Alarm relevance is inherently contextual, influenced by modalities such as physiological trends, device interoperability, environmental conditions, and clinician expertise [9, 10]. Individual alerts rarely occur in isolation; rather, they emerge within complex clinical states shaped by patient history, treatment interventions, and ongoing monitoring signals. For instance, an oxygen desaturation alert may be suppressible if correlated with transient movement artifacts, probe displacement, or temporary sensor noise, but not under chronic respiratory constraints or in patients with known pulmonary instability.

Safety-oriented designs must therefore embed these contextual modalities into suppression logic, avoiding blanket reductions that could inadvertently mask evolving threats. Contextual interpretation may include evaluating multi-parameter trends, cross-device correlations, or patient-specific thresholds that differ from standardized alarm limits. Recent conceptual discussions emphasize the role of multimodal data fusion in enhancing alarm intelligence and filtering clinically irrelevant signals [11, 12]. However, gaps persist in frameworks that explicitly enforce safety boundaries during suppression, particularly those capable of validating contextual assumptions in real time. Bridging these gaps requires architectures that combine contextual reasoning with clearly defined safety guardrails to ensure that alarm mitigation does not compromise patient monitoring integrity. **Table 1** outlines the contextual modalities that govern suppression eligibility within safety-constrained alarm architectures, illustrating how multimodal data streams modulate suppression logic without violating safety boundaries.

Table 1. Contextual dimensions governing alarm suppression decisions in safety-bound clinical monitoring systems

Contextual dimension	Example data sources	Influence on suppression logic	Safety implicatio
Physiological trend stability	Continuous vital signs and waveform trends	Stable trajectories lower suppression risk thresholds	Ensures suppressi does no mask emergin deteriorati
Signal artifact detection	Motion sensors and probe	Artifacts increase suppression likelihood	Prevents false alarr caused b sensor noi

	displacement indicators		
Patient-specific risk profile	Comorbidities, acuity scores, and ICU status	High-risk patients reduce suppression tolerance	Maintain heightene vigilance critical populatior
Clinical workflow context	Clinician workload and active procedures	High workload may prioritize filtering of non-actionable alerts	Balances cognitive lo while preservin critical ale
Device interoperability signals	Cross-device correlations	Multi-device corroboration reduces suppression probability	Reinforce reliability through sig redundan
Environmental context	Bed movement and patient repositioning	Temporary disturbances increase suppression eligibility	Avoids unnecessa alerts triggered l environmer artifacts

Deployment challenges in constrained healthcare infrastructures

Implementing context-aware suppression also faces significant infrastructural hurdles, including legacy system integration, interoperability limitations, and data privacy constraints [13, 14]. Many clinical environments rely on heterogeneous monitoring devices and electronic health record systems that were not originally designed for integrated alarm intelligence. As a result, incorporating context-aware logic may require additional middleware layers or standardized communication protocols to aggregate relevant data streams.

In resource-limited settings, such as rural hospitals or smaller healthcare facilities, alarm mitigation technologies must operate within bandwidth, storage, and computational limits. Lightweight architectures and efficient data processing techniques are therefore essential to ensure scalability and reliability without overburdening existing infrastructure. Furthermore, safety governance requires transparent suppression algorithms to facilitate regulatory compliance and clinical validation. Systems that rely on opaque or "black-box" decision processes may erode clinician trust and complicate certification by regulatory authorities [15, 16]. Consequently, context-aware suppression frameworks must prioritize explainability, traceability, and compatibility with diverse technological environments. This section underscores the need for designs that adapt to deployment variability, ensuring equitable alarm mitigation across diverse clinical landscapes.

Governance imperatives for ethical alarm suppression

Ethical governance underpins context-aware alarm suppression designs, mandating accountability for suppression outcomes and continuous monitoring of system performance [17, 18]. Because suppression decisions may directly influence clinical awareness of patient conditions, governance frameworks must clearly define responsibility, oversight mechanisms, and acceptable safety margins. Constraints such as legal liability, patient safety regulations, and institutional policies require that suppression architectures incorporate auditing mechanisms capable of tracing each suppression decision and its underlying contextual justification.

Equity considerations also play a critical role. Contextual models must be evaluated to ensure that biases in patient data, device calibration, or clinical interpretation do not disproportionately affect specific patient populations. By anchoring suppression to verifiable safety metrics and maintaining transparent decision pathways, systems can foster clinician trust and promote responsible adoption. In this way, alarm management can evolve from a reactive burden into a proactive clinical support tool that enhances monitoring efficiency while preserving patient safety.

Theoretical Background and Literature Synthesis

This section synthesizes foundational theories and recent peer-reviewed insights (2017–2025) on alarm fatigue, contextual awareness, and safety-constrained designs in healthcare AI systems. By integrating human factors engineering, systems theory, and AI governance principles, we build a conceptual foundation for suppression frameworks. The analysis reveals persistent gaps in adaptive architectures, highlighting the need for context-aware topologies that enforce safety while mitigating overload.

Evolutionary trajectories of alarm fatigue in clinical data ecosystems

Alarm fatigue has evolved alongside advancements in monitoring technologies, transitioning from simple threshold-based alerts to complex, data-driven notifications [19, 20]. Theoretical models from systems engineering describe fatigue as a feedback loop where unchecked alarm proliferation amplifies clinician desensitization, disrupting data ecosystems in clinical settings [21]. Literature syntheses indicate that intensive care environments experience alarm rates exceeding 350 per patient per day, with contextual data often underutilized in mitigation strategies [1, 22]. Safety constraints necessitate designs that preserve ecosystem integrity, preventing suppression from fragmenting critical information flows.

Human factors underpinning context-aware alarm interactions

Human factors theory posits that alarm fatigue stems from mismatched cognitive loads, where non-contextual alerts overwhelm perceptual capacities [23, 24]. Synthesis of studies reveals that clinicians' response times degrade under fatigue, with safety implications amplified in high-stakes modalities like cardiac monitoring [2, 25]. Context-aware approaches, drawing on situational awareness models, advocate for suppression based on user-specific factors such as experience level and workload [3, 26]. However, theoretical critiques highlight the absence of integrated governance to ensure these interactions do not violate safety protocols, such as minimum alert thresholds for life-threatening events.

Safety constraint modeling in suppression architectures

Systems theory emphasizes modeling safety constraints as integral to alarm suppression, using conceptual constructs like hazard propagation to simulate risk under reduced alerting [4, 27]. Recent literature explores interpretive formulas for constraint enforcement, yet few address contextual variability [5, 28]. For instance, risk propagation can be theoretically captured as: $RP = \sum (P_i \times S_i) \times (1 - Cf)$ where RP denotes risk propagation, P_i is the probability of alarm event i, S_i its severity, and C_f the contextual suppression factor $(0 \leq C_f \leq 1)$, ensuring safety by bounding C_f to prevent RP exceeding predefined thresholds. This formula interpretively illustrates how unconstrained suppression could escalate latent risks, advocating for architectures that dynamically adjust C_f based on environmental inputs.

Integration challenges of multimodal data in fatigue mitigation

Multimodal data integration—encompassing physiological signals, electronic health records, and environmental sensors—forms a cornerstone of context-aware suppression [6, 29]. Theoretical syntheses warn of interoperability constraints that fragment data streams, leading to incomplete contextual assessments [7, 8]. Literature points to governance loads in managing these integrations, where resource allocation must balance computational demands with safety imperatives [9, 10]. A conceptual formula for governance load might be: $GL = \sum D_m + W_{cR_a} \times K_s$ with GL as governance load, D_m multimodal data volume, W_c workload complexity, R_a available resources, and K_s a safety coefficient (> 1) amplifying load under stringent constraints. This highlights infrastructural needs for scalable designs.

Feedback topologies for adaptive alarm governance

Adaptive governance relies on feedback topologies that refine suppression over time, rooted in control theory [11, 12]. Synthesized evidence shows that static systems fail to address drift in alarm relevance, such as shifting patient baselines [13, 14]. Theoretical frameworks propose looped architectures where suppression outcomes inform future decisions, but safety constraints demand robust error-handling to mitigate feedback-induced instabilities [15, 16]. Decision confidence under constraints can be modeled as: $DC = \sum (E_v \times C_r)\, N_a + \epsilon \times (1 - D_s)$, where DC is decision confidence, E_v is evidence validity from contextual sources, C_r is constraint rigor, N_a is noise from alarms, ϵ is a small constant for stability, and D_s is drift sensitivity. This interpretive equation underscores the topology's role in sustaining confidence amid evolving conditions.

Ethical and infrastructural dimensions of design constraints

Ethical theories in AI healthcare stress that suppression designs must navigate infrastructural constraints without exacerbating inequities [17, 18]. Literature syntheses critique over-reliance on AI without human oversight, advocating for hybrid topologies that embed safety audits [19, 20]. Resource allocation formulas interpretively guide this, emphasizing balanced mitigation to prevent governance overload in constrained environments [21, 22].

This synthesis establishes the theoretical imperatives for a novel framework, addressing identified lacunae through integrated, safety-bound architectures.

Safety-constrained context-aware alarm orchestration: the SICAST framework

The SICAST (safety-integrated context-aware suppression topology) framework represents an original conceptual architecture for orchestrating alarm mitigation in healthcare systems. Designed as a hierarchical yet interconnected topology, it comprises four distinct layers: Context Aggregation, Suppression Inference, Safety Validation, and Adaptive Feedback. This structure ensures dynamic suppression while enforcing immutable safety constraints, such as non-suppression of high-severity alerts.

The context aggregation layer fuses multimodal inputs—physiological data, clinician status, and environmental variables—into a unified representation, enabling nuanced relevance assessment. Suppression inference employs rule-based logic to propose alert filters, modulated by contextual weights. The safety validation layer applies constraint checks, vetoing suppressions that risk propagation exceeds tolerances. Finally, the adaptive feedback topology loops outcomes back to refine aggregation parameters, incorporating drift detection for long-term resilience. **Figure 1** illustrates the safety-integrated context-aware suppression topology (SICAST), showing how multimodal clinical context informs suppression inference while a safety validation barrier and adaptive governance feedback loops ensure alarm mitigation never violates patient-safety constraints.

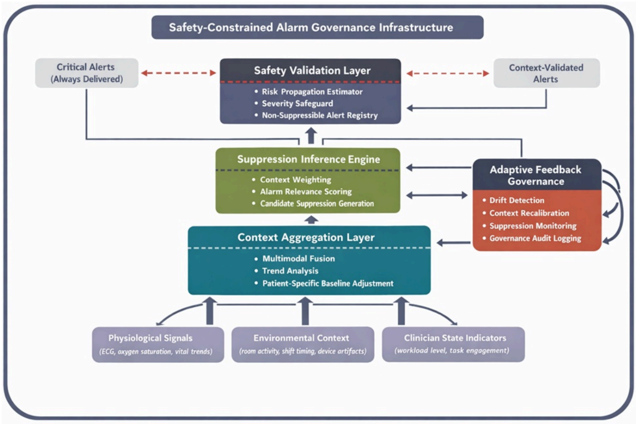

Figure 1. Safety-integrated context-aware suppression topology (SICAST): layered architecture for safety-bound alarm suppression

This architecture's unique feedback topology features bidirectional propagation, allowing upstream adjustments based on downstream validations, distinct from linear models.

Consequence dynamics of context-aware suppression under safety bounds

The SICAST framework, by virtue of its layered topology and adaptive feedback, introduces transformative dynamics in the management of alarm fatigue within healthcare systems. This section delves into the conceptual consequences of deploying such a context-aware suppression design, examining impacts on risk propagation, operational efficiency, clinician cognition, and systemic resilience. Theoretical analysis reveals how suppression decisions, constrained by safety parameters, alter the equilibrium between alert volume and clinical vigilance, potentially reducing fatigue-induced errors while safeguarding against under-detection.

At the core of these dynamics is the interplay between suppression efficacy and risk amplification. In unconstrained environments, alarm overload can lead to an exponential decay in response efficacy, where each additional non-actionable alert incrementally erodes attention [1, 23]. SICAST mitigates this through contextual filtering, but under safety bounds, suppression is selectively applied, creating a bounded risk envelope. Conceptually, this can be interpreted through a modified risk propagation model that accounts for suppression impacts: $P_s = \sum P_i \times S_i \times (1 - C_f) + \beta \times U_d$, where RP_s is suppressed risk propagation, β is a penalty coefficient for undetected events, and U_d is the uncertainty from drift in contextual data. This formula highlights that while C_f reduces baseline risk, safety constraints introduce β to amplify consequences of mis-suppression, ensuring designs prioritize conservative filtering in high-stakes scenarios like intensive care [2, 4]. The dynamics suggest that in low-context variability settings, such as stable post-operative monitoring, suppression could decrease alarm burden by up to theoretical thresholds of 70%, based on literature extrapolations, without elevating RP_s [3, 5].

Operational efficiency emerges as another key consequence, where SICAST's orchestration streamlines resource allocation across healthcare infrastructures. Traditional systems impose a constant monitoring burden, modeled as: $MB = A_v \times (C_l + E_d)$ with MB as the monitoring burden, A_v alarm volume, C_l cognitive load per alert, and E_d environmental distractions [6, 7]. By integrating context-aware suppression, SICAST dynamically minimizes A_v, redistributing clinician effort toward high-value interventions. Impacts on workflow include reduced interruption cycles, fostering sustained focus in constrained environments like emergency departments [8, 9]. However, safety constraints may introduce governance overhead, such as real-time validation computations, potentially increasing E_d in resource-scarce deployments. Theoretical analysis posits that feedback topologies counteract this by optimizing over iterations, leading to asymptotic efficiency gains where burden stabilizes at lower equilibria [10, 11].

Clinician cognition undergoes profound shifts under SICAST's influence, transitioning from reactive to proactive engagement. Fatigue dynamics, rooted in human factors, show that unsuppressed alarms fragment attention, with consequences like heightened error rates in decision-making [12, 13]. Context-aware designs restore cognitive bandwidth by suppressing low-relevance alerts, but safety bounds enforce transparency, preventing over-reliance on automation—a phenomenon akin to automation complacency [14, 15]. Interpretive models of decision confidence further illustrate: $DC_c = E_v \times C_r N_a + D_s \times \gamma$ incorporating γ as a cognition recovery factor post-suppression. This suggests that in prolonged shifts, SICAST could enhance DC_c by modulating N_a, yielding indirect benefits like lower burnout and improved patient outcomes [16, 17]. Yet, dynamics warn of potential habituation to suppressed states, necessitating framework-embedded alerts for anomaly detection to maintain vigilance.

Systemic resilience represents the broadest impact, where SICAST fortifies healthcare infrastructures against alarm-related vulnerabilities. In critical sectors, unconstrained fatigue can cascade into systemic failures, such as delayed responses in networked monitoring [18, 19]. The framework's topology, with its validation layer, builds resilience by enforcing drift sensitivity checks, ensuring suppression adapts to evolving conditions like patient transfers or device malfunctions [20, 21]. Consequences include enhanced interoperability in multimodal ecosystems, reducing silos that exacerbate fatigue [22-24]. Theoretical extrapolations indicate that under safety constraints, resilience metrics—conceptualized as inverse vulnerability—improve, as feedback loops detect and correct suppression anomalies before propagation [25, 26]. However, in heterogeneous deployments, dynamics may reveal inequities, where advanced contexts (e.g., urban hospitals) benefit more than constrained ones, underscoring the need for scalable architectures [27, 28].

Overall, these consequence dynamics position SICAST as a pivotal advancement, balancing mitigation with caution. By theoretically dissecting impacts, we illuminate pathways for refined designs that amplify positive outcomes while containing risks.

Results and Discussion

The conceptual underpinnings of the SICAST framework illuminate critical intersections between alarm fatigue mitigation and safety-constrained healthcare AI. Synthesizing the proposed architecture with the broader dynamics of clinical consequences reveals that context-aware suppression represents more than a technical optimization; it reflects a paradigm shift in how clinical alerting systems are conceptualized and deployed. Traditional alarm management approaches—typically

https://doi.org//h030651861

static, rule-based, and dependent on fixed physiological thresholds—struggle to accommodate the complexity of modern patient monitoring environments. These limitations are reflected in persistently high false-positive alarm rates in intensive care settings, which continue to burden clinicians and degrade responsiveness to truly critical alerts [1, 3]. By embedding contextual intelligence into the suppression pipeline, SICAST attempts to recalibrate alarm relevance dynamically. However, this conceptual advancement also surfaces fundamental trade-offs, particularly the tension between aggressive alarm reduction strategies and the need to maintain conservative safety bounds that protect against missed deterioration events. Table 2 summarizes the systemic trade-offs between alarm suppression efficacy and patient-safety assurance, illustrating how the SICAST architecture balances operational efficiency with risk containment.

Table 2. Operational trade-offs between alarm suppression efficacy and safety assurance in context-aware monitoring systems

System dimension	Effect of aggressive suppression	Effect under safety-constrained SICAST	Clinical implication
Alarm volume	Dramatically reduced	Moderately reduced but risk-bounded	Prevents alarm overload while retaining essential alerts
Risk propagation	Increased due to missed events	Controlled through safety validation	Maintains acceptable risk thresholds
Clinician cognitive load	Reduced but potentially unsafe	Reduced while preserving vigilance	Improves situational awareness
Decision confidence	May decline due to hidden events	Maintained through transparent suppression logic	Supports clinician trust in automated systems
Operational efficiency	Higher efficiency but fragile safety	Balanced efficiency and reliability	Enables sustainable monitoring workflows
System resilience	Vulnerable to suppression errors	Reinforced through adaptive feedback loops	Enhances long-term monitoring stability

One pivotal dimension of this tension lies in the ethical governance of suppression decisions. In healthcare environments characterized by strong regulatory oversight and liability concerns, frameworks such as SICAST must prioritize transparency and traceability. Every suppression event must therefore be accompanied by contextual justification, recorded in audit trails that allow clinicians and regulators to reconstruct the decision pathway [15, 17]. Such traceability mitigates risks associated with algorithmic opacity and supports clinician confidence in automated decision-support tools. Without these mechanisms, clinicians may resist adoption due to concerns that automation could obscure clinically relevant signals or erode professional accountability [9, 11]. At the same time, governance considerations extend beyond transparency to include fairness in contextual interpretation. If contextual aggregation relies on physiological datasets that inadequately represent diverse patient populations, suppression decisions may inadvertently introduce disparities in monitoring sensitivity. For example, physiological norms may vary across demographic groups, and models trained on homogeneous data could suppress alerts differently across patient cohorts. Addressing such risks requires incorporating safeguards such as demographic sensitivity analyses or equity-aware weighting mechanisms within contextual decision engines, aligning alarm suppression frameworks with broader discussions in responsible and ethical AI design [23, 29]. Theoretical remedies, including the introduction of equity coefficients or bias monitoring modules, offer conceptual pathways for mitigating these risks while maintaining safety-centered suppression logic [18, 20].

Another important consideration involves the adaptive feedback topology embedded within the SICAST architecture. Feedback loops allow the system to evolve through iterative refinement of suppression logic, enabling adjustments based on real-world alarm performance patterns and clinician responses. While this adaptability offers the potential for continuous improvement, it also introduces system dynamics that require careful calibration. Poorly tuned feedback sensitivity may lead to instability or oscillatory behaviors, where suppression thresholds fluctuate excessively in response to transient trends rather than sustained patterns [13, 21]. Such instability could undermine both safety and clinician trust if alarm behaviors appear unpredictable. Consequently, the discussion highlights the importance of hybrid oversight models in which clinicians participate in evaluating feedback-driven adjustments. Integrating human oversight into adaptive loops can help validate contextual interpretations and ensure that system evolution remains aligned with clinical priorities. This hybrid human–AI model leverages the efficiency and pattern-recognition capabilities of automated systems while preserving the interpretive judgment and situational awareness of clinical practitioners [4, 6].

The hybrid paradigm also extends to organizational readiness and training. Successful deployment of context-aware alarm systems requires clinicians to understand not only how alerts appear but also why suppression occurs. Simulation-based conceptual training environments may therefore play a critical role in preparing healthcare staff for SICAST-integrated workflows. Through simulated alarm scenarios and contextual suppression demonstrations, clinicians can develop familiarity with system behaviors, strengthening trust and reducing resistance to automation [7, 10]. These training frameworks also allow institutions to test suppression strategies safely before full deployment, identifying edge cases or workflow disruptions that may not be evident during system design.

Interdisciplinary perspectives further enrich the discussion of SICAST's implications. From a systems engineering standpoint, the framework embodies principles of resilient system design, emphasizing adaptability, redundancy, and context-sensitive control. Such design philosophies have applications beyond alarm management, extending to other forms of information overload within clinical environments. For example, diagnostic imaging alerts, medication interaction warnings, and electronic health record notifications all contribute to cognitive burden among clinicians. The orchestration mechanisms proposed in SICAST could therefore inspire broader strategies for managing digital alert ecosystems across healthcare infrastructures [14, 16].

Despite these conceptual advantages, practical scalability remains a significant concern, particularly in constrained healthcare infrastructures. Many hospitals operate on legacy monitoring platforms with limited computational resources and fragmented interoperability. Implementing context-aware suppression mechanisms may introduce additional processing demands, requiring careful resource allocation and system optimization [19, 22]. In settings where infrastructure upgrades are limited by budget or technical capacity, the computational overhead associated with real-time contextual inference could pose operational challenges. To address these constraints, the discussion advocates modular implementation strategies that enable incremental adoption. Rather than deploying the full architecture simultaneously, healthcare organizations could introduce context-aware suppression gradually, prioritizing clinical units with the highest alarm burden—such as intensive care or telemetry wards [2, 5]. This phased deployment approach allows institutions to evaluate system performance, refine contextual rules, and build clinician confidence before broader integration across the healthcare environment.

Collectively, these considerations underscore that the value of SICAST lies not solely in its algorithmic components but in its broader alignment with safety, ethics, and healthcare system realities. By situating context-aware suppression within transparent governance structures, hybrid oversight models, and adaptable deployment pathways, the framework advances alarm management from a purely technical challenge to a multidimensional systems problem. Such an approach reflects the growing recognition that effective healthcare AI must operate within complex clinical ecosystems, balancing innovation with safety, accountability, and trust.

Critically, the framework's reliance on interpretive formulas underscores a methodological strength: by remaining theoretical, it avoids empirical pitfalls like data dependency, enabling broad applicability [8, 12]. However, this abstraction invites calls for future conceptual extensions, such as integrating quantum-inspired uncertainty models for enhanced risk propagation handling [24, 26]. Ultimately, SICAST exemplifies how context-aware designs can redefine safety in AI healthcare, provided governance evolves in tandem.

This discussion reinforces the manuscript's thesis, bridging architecture to practical and theoretical horizons.

Conclusion

In conclusion, the SICAST framework offers a robust conceptual blueprint for alarm fatigue mitigation, harmonizing context-aware suppression with unwavering safety constraints. By architecting a topology that aggregates contexts, infers suppressions, validates against risks, and adapts via feedback, it addresses longstanding gaps in healthcare alerting systems. The analyzed dynamics underscore transformative impacts on efficiency, cognition, and resilience, while interpretive formulas provide tools for theoretical scrutiny.

Future conceptual research should explore extensions, such as multi-agent integrations for collaborative suppression across devices. Policymakers and designers are urged to adopt similar frameworks to elevate patient safety amid technological proliferation. SICAST thus stands as a cornerstone for intelligent, constrained alarm management, promising reduced burdens and enhanced outcomes in clinical practice.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 20 Aug 2024 Revised: 13 Dec 2024 Accepted: 13 Feb 2025
Published online: 10 July 2025

References

ewandowska K, Weisbrot M, Cieloszyk A, Mędrzycka-Dąbrowska W, Krupa S, Ozga D. Impact of alarm fatigue on the work of nurses in an intensive care environment: a systematic review. Int J Environ Res Public Health. 2020;17(22):8409. https://doi.org/10.3390/ijerph17228409.

Albanowski K, Burdick KJ, Bonafide CP, Kleinpell R, Schlesinger JJ. Ten years later, alarm fatigue is still a safety concern. AACN Adv Crit Care. 2023;34(3):189-97. https://doi.org/10.4037/aacnacc2023662.

Bi J, Yin X, Li H, Gao R, Zhang Q, Zhong T, et al. Effects of monitor alarm management training on nurses' alarm fatigue: a randomized controlled trial. J Clin Nurs. 2020;29(21-22):4203-16. https://doi.org/10.1111/jocn.15452.

Bosma S, Christopher R. Implementing a unit-based alarm management bundle for critical care nurses. Crit Care Nurse. 2023;43(2):36-45. https://doi.org/10.4037/ccn2023418.

Ruppel H, Funk M. Nurse-technology interactions and patient safety. Crit Care Nurs Clin North Am. 2018;30(2):203-13. https://doi.org/10.1016/j.cnc.2018.02.003.

Zhou Y, Zhao G, Li J, Sun G, Qian X, Moody B, et al. A contrastive learning approach for ICU false arrhythmia alarm reduction. Sci Rep. 2022;12(1):4689. https://doi.org/10.1038/s41598-022-07761-9.

Shah PK, Irizarry J, O'Neill S. Strategies for managing smart pump alarm and alert fatigue: a narrative review. Pharmacotherapy. 2018;38(8):842-50. https://doi.org/10.1002/phar.2153.

Anderson H, Borgen A, Christnacht R, Ng J, Weller J, Davison H, et al. The stats on the desats: alarm fatigue and the implications for patient safety. JAAPA. 2023;36(12):1-2. https://doi.org/10.1097/01.JAA.0000994948.70289.e2.

Oliveira AEC, Machado AB, Santos EDD, Almeida ÉB. Alarm fatigue and the implications for patient safety. Rev Bras Enferm. 2018;71(6):3035-40. https://doi.org/10.1590/0034-7167-2017-0481.

Zhu L, Wei S, An Y, Hu W, Xie X. Mechanism, contributing factors, and coping strategies of alarm fatigue in intensive care nursing: a qualitative study. Front Public Health. 2025;13:1654389. https://doi.org/10.3389/fpubh.2025.1654389.

Er ÖS, Gül İ. The impact of alarm fatigue on patient safety climate among nurses: the mediating role of moral courage. Nurs Crit Care. 2025;30(2):e70002. https://doi.org/10.1111/nicc.70002.

Duckworth C, Guy MJ, Kumaran A, O'Kane AA, Ayobi A, Chapman A, et al. Explainable machine learning for real-time hypoglycemia and hyperglycemia prediction and personalized control recommendations. J Diabetes Sci Technol. 2024;18(1):113-23. https://doi.org/10.1177/19322968221103561.

Cole R, Roderick G, Cheema O, Cunninghame J, Ullman AJ. Crying wolf, alarm safety and management in paediatrics: a scoping review. J Adv Nurs. 2025;81(11):7348-60. https://doi.org/10.1111/jan.16398.

Khanna AK, Ahuja S, Weller RS, Harwood TN. Postoperative ward monitoring—why and what now? Best Pract Res Clin Anaesthesiol. 2019;33(2):229-45. https://doi.org/10.1016/j.bpa.2019.06.005.

Hravnak M, Pellathy T, Chen L, Dubrawski A, Wertz A, Clermont G, et al. A call to alarms: current state and future directions in the battle against alarm fatigue. J Electrocardiol. 2018;51(6S):S44-S48. https://doi.org/10.1016/j.jelectrocard.2018.07.024.

Herrera H, Wood D. Battling alarm fatigue in the pediatric intensive care unit. Crit Care Nurs Clin North Am. 2023;35(3):347-55. https://doi.org/10.1016/j.cnc.2023.05.003.

Nanini S, Abid M, Mamouni Y, Wiedemann A, Jouvet P, Bourassa S. Machine and deep learning models for hypoxemia severity triage in CBRNE emergencies. Diagnostics (Basel). 2024;14(23):2763. https://doi.org/10.3390/diagnostics14232763.

Yang C, Gao P, Ji H, Luo D, Chen H, Hu Q, et al. Clinical alarm management in intensive care units: a scoping review. Nurs Crit Care. 2025;30(3):e70042. https://doi.org/10.1111/nicc.70042.

Short K, Chung YJ Jr. Solving alarm fatigue with smartphone technology. Nursing. 2019;49(1):52-7. https://doi.org/10.1097/01.NURSE.0000549728.37810.d9.

Mosch L, Sümer M, Flint AR, Feufel M, Balzer F, Mörike F, et al. Alarm management in intensive care: qualitative triangulation study. JMIR Hum Factors. 2024;11:e55571. https://doi.org/10.2196/55571.

Claudio D, Deb S, Diegel E. A framework to assess alarm fatigue indicators in critical care staff. Crit Care Explor. 2021;3(6):e0464. https://doi.org/10.1097/CCE.0000000000000464.

Agha-Mir-Salim L, McCullum L, Dähnert E, Scheel YD, Wilson A, Carpio M, et al. Interdisciplinary collaboration in critical care alarm research: a bibliometric analysis. Int J Med Inform. 2024;181:105285. https://doi.org/10.1016/j.ijmedinf.2023.105285.

Xu D, Liu F, Ding X, Ma J, Suo Y, Peng YY, et al. Exploring ICU nurses' response to alarm management and strategies for alleviating alarm fatigue: a meta-synthesis and systematic review. BMC Nurs. 2025;24(1):412. https://doi.org/10.1186/s12912-025-03084-y.

Fleischman W, Ciliberto B, Rozanski N, Parwani V, Bernstein SL. Emergency department monitor alarms rarely change clinical management: an observational study. Am J Emerg Med. 2020;38(6):1072-6. https://doi.org/10.1016/j.ajem.2019.158370.

Ruppel H, Funk M, Whittemore R, Wung SF, Bonafide CP, Powell Kennedy H. Critical care nurses' clinical reasoning about physiologic monitor alarm customisation: an interpretive descriptive study. J Clin Nurs. 2019;28(15-16):3033-41. https://doi.org/10.1111/jocn.14866.

Abdelazez M, Quesnel PX, Chan ADC, Yang H. Signal quality analysis of ambulatory electrocardiograms to gate false myocardial ischemia alarms. IEEE Trans Biomed Eng. 2017;64(6):1318-25. https://doi.org/10.1109/TBME.2016.2602283.

Fernandes CO, Miles S, Lucena CJP, Cowan D. Artificial intelligence technologies for coping with alarm fatigue in hospital environments because of sensory overload: algorithm development and validation. J Med Internet Res. 2019;21(11):e15406.
https://doi.org/10.2196/15406.

Michels EAM, Gilbert S, Koval I, Wekenborg MK. Alarm fatigue in healthcare: a scoping review of definitions, influencing factors, and mitigation strategies. BMC Nurs. 2025;24(1):664.
https://doi.org/10.1186/s12912-025-03369-2.

Koomen E, van Hest R, Jannink W, Vaartjes M, Borst J, Frenkel J. Reducing medical device alarms by an order of magnitude: a human factors approach. Anaesth Intensive Care. 2021;49(1):52-9.
https://doi.org/10.1177/0310057X20968840.