Digital Health and Medical Intelligence

https://dhmijournal.com

Forecasting Alert Fatigue in Emergency Department Workflows with Algorithmic Risk Scores

Daniel E. Allen*, Daniel Price

Graduate School of Biomedical Sciences, Icahn School of Medicine at Mount Sinai, New York, USA

Corresponding author: daniel.e.allen@gmail.com

Abstract

The integration of algorithmic clinical decision support systems in emergency departments has fundamentally transformed patient care by providing real-time risk scores and automated warnings. However, the proliferation of these notifications has precipitated a critical unintended consequence known as alert fatigue, wherein clinicians become desensitized to warnings, leading to increased override rates and potential patient harm. This paper investigates the predictability of alert fatigue derived from algorithmic risk scores through a comprehensive mixed evaluation methodology situated within the complex workflow of a high-acuity emergency department. By synthesizing quantitative data from electronic health record access logs with qualitative insights gathered from clinician shadowing and structured interviews, this research delineates the multifaceted nature of alert fatigue. We propose a framework for understanding how continuous exposure to varying tiers of algorithmic risk scores deteriorates clinical responsiveness. The study identifies key workflow disruptions, cognitive burden indicators, and behavioral adaptations that emergency personnel develop in response to high-frequency alerts. The findings reveal that alert fatigue can be reliably predicted not only by the sheer volume of alerts but by the contextual misalignment between algorithmic risk scores and the dynamic clinical reality of the emergency department. By elucidating the underlying mechanisms of cognitive overload, this paper provides critical insights for the redesign of clinical decision support systems, advocating for context-aware, threshold-calibrated alert mechanisms that prioritize clinician cognitive bandwidth and ultimately safeguard patient safety.

Keywords: Alert Fatigue; Emergency Department; Risk Scoring Algorithms; Mixed Methods Evaluation; Algorithmic Risk Scores

1. Introduction

1.1 Context and Problem Statement

The modern emergency department is a highly dynamic, time-sensitive environment characterized by continuous influxes of critically ill patients, rapid triage requirements, and immense cognitive demands placed upon healthcare providers. To assist clinicians in managing this complexity, healthcare institutions have widely adopted sophisticated clinical decision support systems embedded within electronic health records. These systems leverage complex algorithms to generate risk scores and automated alerts intended to prevent medication errors, identify deteriorating clinical trajectories, and ensure compliance with best practice guidelines [1]. However, the aggressive proliferation of these automated warnings has generated an overwhelming volume of digital noise. Emergency department clinicians are frequently bombarded with hundreds of notifications per shift, a substantial proportion of which are deemed clinically insignificant or overly cautious for the specific context of emergency care. This incessant barrage of notifications inevitably leads to alert fatigue, a well-documented psychological and behavioral phenomenon where human operators become increasingly desensitized to warnings, resulting in prolonged response times, elevated rates of alert overrides, and the dangerous potential for ignoring genuinely critical alerts [2].

The fundamental problem lies in the disconnect between the algorithmic generation of risk scores and the nuanced, context-dependent reality of emergency department workflows. Algorithms typically operate on predefined physiological thresholds or standardized pharmacological databases, generating alerts without an adequate understanding of the immediate clinical picture or the specific phase of patient care [3]. Consequently, a clinician might receive a severe sepsis warning for a patient who has already been diagnosed and is currently receiving appropriate broad-spectrum intravenous antibiotics, rendering the alert redundant and disruptive. The sheer volume of such false positive or clinically irrelevant alerts significantly erodes the perceived utility of the clinical decision support system, transforming a purported safety net into a substantial source of cognitive friction and occupational burnout [4]. Predicting when and how this fatigue sets in remains a complex challenge, as it requires moving beyond simple counting of alerts to understanding the contextual, algorithmic, and human factors that converge in the clinical workflow.

1.2 Research Objectives and Scope

This paper aims to bridge the existing knowledge gap by developing a comprehensive understanding of how algorithmic risk scores contribute to and predict the onset of alert fatigue within emergency department workflows. The primary objective is to investigate the predictability of clinician desensitization by analyzing the interplay between quantitative alert characteristics such as risk score severity, frequency, and display modality, and qualitative workflow factors such as clinical task interruption and cognitive workload. By employing a mixed evaluation methodology, this research seeks to capture both the measurable behavioral responses to alerts and the subjective experiential impact on the healthcare provider.

The scope of this investigation encompasses the analysis of retrospective electronic health record data, specifically focusing on alert generation logs, user response times, and override justifications, coupled with prospective qualitative data derived from direct observational shadowing and structured interviews with emergency department personnel. By triangulating these distinct data sources, the study endeavors to formulate a robust, multidimensional model of alert fatigue progression. This model will not only identify the critical thresholds at which algorithmic risk scores lose their efficacy but will also highlight the contextual modifiers that exacerbate or mitigate cognitive overload. Ultimately, the findings are intended to inform the strategic redesign of clinical decision support architectures, advocating for the implementation of dynamic, context-aware algorithms that respect the cognitive limitations of human operators and align seamlessly with the high-stakes workflow of emergency medicine.

2. Background and Literature Review

2.1 Evolution of Clinical Decision Support Systems

The integration of automated assistance into clinical practice has a rich history, evolving from rudimentary rule-based drug interaction checkers to the contemporary landscape of predictive analytics and machine learning-driven risk stratification tools. Early iterations of these systems were designed primarily to intercept basic prescribing errors, relying on static databases and deterministic logic [5]. While these foundational systems demonstrated initial success in reducing adverse drug events, their rigid architecture often failed to account for clinical nuance, resulting in a high volume of nuisance alerts. As electronic health records became ubiquitous, the scope of clinical decision support expanded exponentially. Modern systems now continuously ingest vast streams of physiological data, laboratory results, and demographic information, applying sophisticated algorithms to generate real-time risk scores for conditions such as sepsis, acute kidney injury, and clinical deterioration [6].

Despite these technological advancements, the fundamental interaction paradigm between the system and the clinician remains largely unchanged. The primary mechanism of communication is still the interruptive alert, a modal pop-up or prominent visual cue that demands immediate attention and cognitive processing from the user [7]. This reliance on interruptive notification paradigms, inherited from earlier, simpler systems, has proven fundamentally incompatible with the complex, multitasking environment of the emergency department. The persistent failure to adapt the delivery mechanism to the sophistication of the underlying algorithms has exacerbated the tension between automated safety protocols and human cognitive capacity, creating an environment ripe for the development of severe alert fatigue [8].

2.2 The Psychology and Measurement of Alert Fatigue

Alert fatigue is deeply rooted in human psychological responses to chronic sensory and cognitive stimulation. Drawing upon principles of sensory adaptation and signal detection theory, researchers have established that continuous exposure to high-frequency, low-fidelity warnings predictably degrades an operator's ability to discriminate between genuine threats and false alarms [9]. In the clinical domain, this desensitization manifests as a subconscious recalibration of risk perception, where the default assumption shifts from viewing an alert as a critical warning to perceiving it as a likely nuisance. This psychological shift is accompanied by measurable behavioral changes, most notably a drastic increase in the rate at which alerts are rapidly dismissed or overridden without thorough clinical evaluation.

Previous attempts to quantify alert fatigue have predominantly relied on retrospective analyses of electronic health record logs, calculating basic metrics such as alert volume per shift, average time to override, and the raw percentage of alerts ignored. While these quantitative measures provide a macroscopic view of the problem, they frequently fail to capture the contextual drivers of the fatigue [10]. For instance, a high override rate might not solely indicate desensitization but could reflect a fundamentally flawed algorithm that consistently generates clinically inappropriate recommendations for a specific patient population. Furthermore, pure log analysis cannot assess the cumulative cognitive toll that repeated workflow interruptions exact on the clinician. Recent literature emphasizes the necessity of incorporating qualitative assessments, such as cognitive workload inventories and workflow mapping, to fully understand the multidimensional impact of alert fatigue [11]. The current study builds upon this paradigm, utilizing a mixed evaluation approach to synthesize objective behavioral data with subjective clinical experiences, thereby providing a more nuanced and comprehensive understanding of how algorithmic risk scores influence human performance in emergency settings [12].

3. Theoretical Framework for Alert Fatigue

3.1 Conceptualizing the Algorithmic Cognitive Load

To systematically investigate the phenomenon of alert fatigue, this study constructs a theoretical framework that conceptualizes the interaction between algorithmic risk scores and clinician cognition as a dynamic stress-strain relationship. Central to this framework is the concept of cognitive bandwidth, which posits that healthcare providers possess a finite capacity for processing information and making decisions within a given timeframe. When a clinical decision support system generates an alert, it acts as an external demand on this limited bandwidth. The magnitude of this cognitive demand is determined not only by the urgency conveyed by the algorithmic risk score but also by the immediate contextual factors surrounding the clinician at the exact moment of interruption [13].

The framework proposes that alert fatigue is not merely a linear consequence of volume, but a complex outcome mediated by the perceived signal-to-noise ratio of the system. Every time an algorithm generates a risk score that prompts an alert, the clinician engages in a rapid, often subconscious, process of cognitive appraisal. They must evaluate the validity of the algorithm's output against their own clinical assessment of the patient. If the algorithm consistently aligns with clinical reality, trust is maintained, and the cognitive load of processing the alert is relatively low. However, when the system frequently produces false positives or contextually inappropriate warnings, the cognitive cost of appraising and dismissing these alerts accumulates [14]. This chronic cognitive taxation inevitably leads to a state of depletion, where the clinician adopts heuristic processing strategies, such as automated overriding, to conserve mental resources and maintain workflow momentum [15].

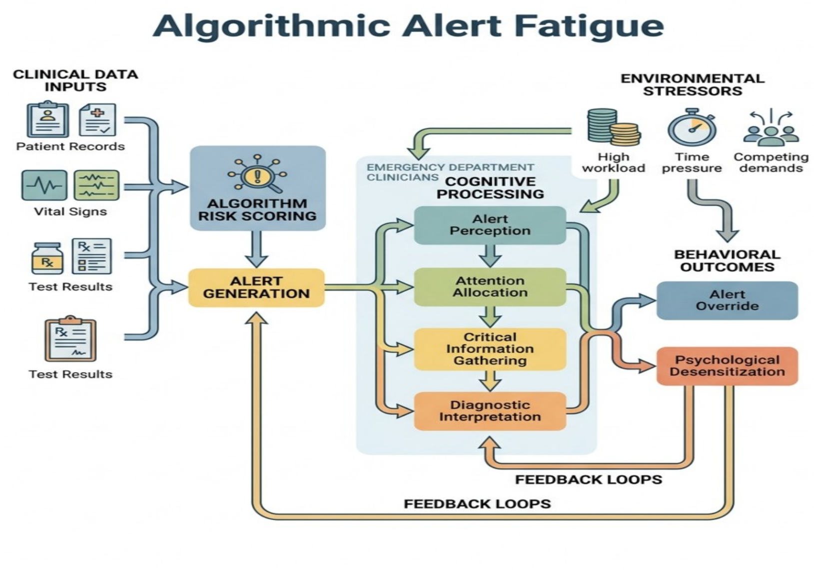

Figure 1: Comprehensive Conceptual Model of Algorithmic Alert Fatigue Progression

3.2 The Mixed Evaluation Paradigm

The theoretical foundation of this research asserts that a single methodological approach is insufficient to capture the intricate dynamics of alert fatigue. Relying solely on quantitative data risks producing a sterile, decontextualized understanding of user behavior, while relying purely on qualitative observations may fail to identify statistically significant patterns operating at scale. Therefore, this study adopts a robust mixed evaluation paradigm, deliberately designing quantitative and qualitative data streams to be complementary and mutually illuminating.

In this integrated model, the quantitative component serves to identify the macro-level behavioral patterns, highlighting which specific algorithmic risk scores and alert types are most frequently overridden and at what point during a shift the override velocity increases. These objective metrics act as diagnostic markers, pinpointing the precise locations within the electronic health record ecosystem where the human-computer interaction is failing. Subsequently, the qualitative component acts as a microscopic lens, providing the contextual narrative necessary to interpret the quantitative anomalies. By engaging directly with the clinicians operating at the interface of these systems, the mixed evaluation paradigm seeks to uncover the rationale behind the override statistics, exploring how factors such as time pressure, competing clinical priorities, and fundamental mistrust of the algorithm drive the measurable manifestations of alert fatigue.

4. Study Design and Methodology

4.1 Research Setting and Participant Selection

This study was conducted within the primary adult emergency department of a large, urban tertiary care teaching hospital. This setting was intentionally selected due to its high patient volume, broad acuity spectrum, and the dense deployment of algorithmic clinical decision support systems integrated within its enterprise electronic health record infrastructure. The emergency department utilizes a diverse array of automated risk scoring algorithms, ranging from basic medication allergy cross-checks to complex predictive models for clinical deterioration and severe sepsis. This environment provides an optimal natural laboratory for observing the real-world impact of high-frequency alerts on clinical workflows.

The study population comprised attending emergency medicine physicians, resident physicians, and emergency department nurses. A purposive sampling strategy was employed to ensure a representative cross-section of clinical experience levels, shift patterns, and primary care roles. Participation in the prospective observational and interview phases was strictly voluntary, and informed written consent was obtained from all subjects prior to any data collection. To minimize the Hawthorne effect during the observational phases, the primary purpose of the shadowing was framed broadly as an evaluation of clinical workflow efficiency rather than a specific investigation into alert overriding behavior. The institutional review board thoroughly reviewed and approved all protocols, ensuring the strict preservation of patient and clinician confidentiality throughout the research process.

4.2 Data Collection Procedures

The data collection architecture was bifurcated into distinct quantitative and qualitative streams, executed concurrently over a continuous three-month study period. The quantitative data stream involved the systematic extraction of backend audit logs from the electronic health record system. This extraction targeted all discrete alert events triggered by algorithmic risk scores during the study timeframe. The captured data elements included the unique identifier of the clinician receiving the alert, the exact timestamp of the notification, the specific algorithm generating the risk score, the severity tier assigned by the system, the clinical context of the alert, and the final action taken by the user, categorized broadly as acceptance, modification, or override. Furthermore, the duration of time elapsed between the presentation of the alert and the user's recorded action was logged to serve as a proxy metric for cognitive engagement.

Parallel to the retrospective data extraction, the qualitative data stream utilized a combination of direct structured observations and semi-structured follow-up interviews. Trained clinical informatics researchers conducted extensive shadowing sessions, observing clinicians across varying shift times, including high-volume daytime hours and resource-constrained night shifts. The observers utilized a standardized data collection tool to document instances of alert generation, specifically noting the clinician's current task, the degree of workflow interruption, visible signs of frustration, and real-time verbalized rationales for overriding notifications. Following the shadowing sessions, the researchers conducted targeted interviews with the observed clinicians. These interviews were structured around critical incident recall, utilizing specific examples of overridden alerts captured during the observation to probe the clinician's underlying decision-making process, their general trust in the algorithmic risk scores, and their subjective experience of cognitive overload [16].

5. Data Processing and Analytical Approach

5.1 Quantitative Log Analysis

The raw data extracted from the electronic health record generated a massive dataset comprising hundreds of thousands of individual alert interactions. The initial phase of data processing involved rigorous cleaning and standardization procedures to remove duplicate entries, system-generated administrative alerts, and incomplete records. Once the dataset was sanitized, the dependent variables for the quantitative analysis were operationalized. The primary metric of interest was the alert override rate, defined as the proportion of total alerts presented to a user that were subsequently dismissed without any corresponding change in clinical ordering behavior. A secondary metric was the response latency, measured in milliseconds, representing the time taken to clear the alert dialog box.

To model the predictors of alert fatigue, a multivariable logistic regression framework was constructed. The dependent variable was the binary outcome of an alert being overridden. The independent predictor variables included characteristics of the alert itself, such as the algorithmic risk score severity, the clinical domain of the alert, and the cumulative number of alerts the specific user had received thus far during their current shift. Additional independent variables encompassed contextual factors, including the hour of the shift, the current patient census in the emergency department, and the clinical role of the user. This analytical approach permitted the isolation of independent risk factors for overriding behavior, allowing the research team to quantify the independent contribution of alert volume and algorithmic complexity to the onset of desensitization [17].

Table 1: Participant Demographics and Professional Characteristics

Characteristic	Total Participants	Physicians	Nurses	Percentage of Cohort
Total Enrollment	120	50	70	100.0
Experience < 5 Years	45	20	25	37.5
Experience 5-10 Years	40	15	25	33.3
Experience > 10 Years	35	15	20	29.2
Day Shift Primary	65	28	37	54.2
Night Shift Primary	55	22	33	45.8

5.2 Qualitative Thematic Synthesis

The qualitative data analysis employed a rigorous thematic synthesis approach to extract meaningful patterns from the observational field notes and interview transcripts. All audio-recorded interviews were transcribed verbatim and subsequently uploaded into a qualitative data analysis software platform. The analytical process began with open coding, where two independent researchers systematically reviewed the transcripts line-by-line, attaching descriptive labels to segments of text that pertained to the clinician's interaction with the clinical decision support system, their perception of algorithmic accuracy, and their experience of workflow disruption [18].

Following the initial open coding phase, the researchers engaged in a collaborative process of axial coding, where the descriptive labels were grouped into broader, overarching thematic categories. This iterative process involved constant comparison between the emerging themes and the raw data to ensure accuracy and fidelity to the clinicians' expressed experiences. Any discrepancies in coding or thematic categorization between the two primary analysts were resolved through moderated discussion and consensus-building with a senior clinical informatics investigator. The final thematic structure provided a comprehensive narrative framework that detailed the psychological progression of alert fatigue, moving from initial frustration with poorly calibrated algorithmic risk scores to the ultimate development of maladaptive behavioral coping mechanisms.

6. Quantitative Results and System Evaluation

6.1 Predictors of Overriding Behavior

The quantitative analysis of the electronic health record logs revealed a staggering volume of algorithmic interruptions within the emergency department workflow. Over the three-month study period, the system generated over eighty thousand distinct alerts for the monitored cohort. The aggregate data demonstrated a profound level of desensitization, with an overall alert override rate exceeding eighty-five percent across all clinical domains. However, the multivariable logistic regression analysis provided critical nuances to this macroscopic figure, identifying several highly significant predictors of overriding behavior.

The most potent independent predictor of an alert being overridden was the cumulative volume of notifications received by the individual clinician during their current shift. The data demonstrated a clear, non-linear deterioration in alert acceptance. During the initial two hours of a shift, the override rate remained relatively stable; however, beyond the fourth hour, the probability of any given alert being summarily dismissed increased exponentially, controlling for the severity of the algorithmic risk score. This finding provides objective empirical evidence for the progressive depletion of cognitive bandwidth. Furthermore, the specific clinical domain of the algorithm significantly influenced the outcome. Algorithmic risk scores related to general medication interaction warnings were overridden at a rate approaching ninety-five percent, whereas alerts derived from predictive models for acute physiological deterioration demonstrated a significantly lower, though still problematic, override rate of approximately sixty percent.

Table 2: Summary of Quantitative Outcomes by Alert Risk Tier

Algorithmic Risk Tier	Total Alerts Generated	Override Rate (%)	Average Response Time (ms)	Immediate Action Taken (%)
Low Acuity (Informational)	35,420	94.2	850	1.5
Moderate Acuity (Warning)	28,150	88.7	1120	4.8
High Acuity (Critical Risk)	12,300	62.4	2450	25.6
Sepsis Predictive Model	4,500	58.1	3100	31.2

6.2 Latency and Contextual Modifiers

Analysis of the response latency data provided further insights into the behavioral manifestations of alert fatigue. For the vast majority of low-acuity and moderate-acuity alerts, the average time to override was significantly less than one second, indicating that clinicians were dismissing the notifications through automated motor responses without dedicating meaningful cognitive attention to the content of the warning. This ultra-rapid dismissal suggests that the interface design of the clinical decision support system has inadvertently trained users to treat alerts as physical obstacles to be cleared rather than vital pieces of clinical information to be synthesized [19].

The quantitative data also highlighted the critical role of environmental context in modifying the impact of algorithmic risk scores. The probability of an alert being overridden demonstrated a strong positive correlation with the current patient census and the aggregate acuity level of the emergency department. During periods of severe overcrowding, the override rate for even high-acuity alerts increased significantly. This statistical reality underscores the limitation of evaluating algorithms in isolation; an algorithmic risk score that appears highly sensitive and specific during in-silico validation may become practically useless when deployed into the chaotic, resource-constrained reality of a department operating at maximum capacity [20].

7. Qualitative Findings and Workflow Integration

7.1 The Erosion of Algorithmic Trust

The thematic analysis of the qualitative interviews provided profound explanatory depth to the high override rates documented in the log analysis. The most pervasive theme to emerge from the clinician narratives was a fundamental erosion of trust in the validity of the algorithmic risk scores. Clinicians repeatedly described a phenomenon akin to "the boy who cried wolf," wherein the high volume of clinically irrelevant alerts systematically destroyed the credibility of the entire clinical decision support system. Many physicians noted that the algorithms seemed inherently risk-averse, calibrated to prioritize maximum sensitivity at the complete expense of specificity. While this design choice might make theoretical sense from a pure liability perspective, the practical consequence in the emergency department was a deluge of false positives that clinicians viewed as actively detrimental to their cognitive focus.

Interviewees frequently highlighted the stark mismatch between the static, generalized logic of the algorithms and the dynamic, highly individualized nature of emergency care. For example, a senior resident expressed deep frustration with continuous alerts warning of renal toxicity for a critically ill patient requiring a necessary contrast-enhanced diagnostic scan, noting that the algorithm possessed no capacity to weigh the immediate life-threatening risk of a missed pulmonary embolism against the theoretical long-term risk of kidney injury. This lack of contextual awareness led clinicians to view the system not as a supportive colleague, but as an adversarial compliance monitor that they had to constantly battle to execute their clinical duties efficiently [21].

7.2 Cognitive Burden and Workflow Fragmentation

The qualitative observations also vividly illustrated the structural damage inflicted upon clinical workflows by persistent algorithmic interruptions. Emergency medicine inherently requires clinicians to maintain a fragile mental model of multiple complex patients simultaneously. Shadowing sessions revealed that interruptive alerts frequently shattered this cognitive equilibrium. Observers documented numerous instances where a physician, deeply engaged in complex medical decision-making or sensitive patient communication, was jarringly interrupted by a low-acuity pop-up notification. Even when the physical act of dismissing the alert took mere seconds, the cognitive cost of task-switching, processing the interruption, and attempting to re-establish the previous train of thought was substantial and cumulative.

Nurses, who interact with the electronic health record with high frequency during medication administration, reported significant psychological distress related to the fear of missing a genuinely critical warning amidst the noise. The necessity to rapidly clear dozens of nuisance alerts merely to administer a standard antibiotic created a dangerous environment where vital information could easily be overlooked through ingrained muscle memory. The qualitative data strongly suggests that predicting alert fatigue requires looking beyond the statistical performance of the algorithmic risk score to understand its disruptive potential within the physical and temporal constraints of the clinical workspace.

8. Conclusion and Future Directions

8.1 Synthesis of Key Findings

This study demonstrates that the phenomenon of alert fatigue in the emergency department is a highly complex, predictable outcome driven by the fundamental misalignment between algorithmic risk generation and the realities of human cognitive processing. The mixed evaluation methodology effectively triangulated the overwhelming volume and rapid dismissal rates captured in the quantitative logs with the subjective experiences of frustration, cognitive depletion, and eroded trust articulated in the qualitative narratives. The findings confirm that while algorithmic risk scores possess immense theoretical potential to enhance patient safety, their current implementation strategy—relying heavily on interruptive, context-blind notifications—is structurally flawed.

The data unequivocally shows that alert override behavior is not a reflection of clinical negligence, but rather a rational, necessary adaptation by healthcare providers attempting to manage an impossible informational load. The probability of alert fatigue can be reliably predicted by analyzing the cumulative volume of notifications, the specificity of the underlying algorithm, and the environmental stress of the clinical setting. When algorithms fail to account for the immediate clinical context, they transform from tools of decision support into instruments of cognitive taxation, ultimately degrading the quality of care they were ostensibly designed to protect.

8.2 Implications and Recommendations

Addressing the crisis of alert fatigue necessitates a fundamental paradigm shift in how clinical decision support systems are conceptualized and deployed. Future engineering efforts must move away from static, threshold-based alerts toward dynamic, machine-learning-driven architectures capable of contextual awareness. Algorithms must be redesigned to incorporate not just physiological data, but workflow state, evaluating the probability that an alert will actually alter clinical management at that precise moment in time.

Furthermore, the mechanisms of information delivery must be diversified. The ubiquitous modal pop-up must be reserved strictly for immediate, life-threatening interventions. Lesser warnings should be integrated seamlessly into the clinical dashboard using non-interruptive visual cues, allowing the clinician to access the algorithmic insights at their own cognitive convenience [22]. Ultimately, the success of algorithmic risk scores in the emergency department depends not on the mathematical elegance of the model, but on its empathetic integration into the demanding, high-stakes workflow of the human clinicians it is intended to serve. Future research must prioritize robust, in-situ usability testing and continuous, closed-loop feedback mechanisms to ensure that clinical decision support systems evolve into true partners in patient care.

References

1. Yoon, W.J. Endoscopic Transaxillary Augmentation Mammoplasty; Springer: Singapore, 2019.

2. Silicone Breast Implant Modification Review: Overcoming Capsular Contracture | Biomaterials Research. Available online: https://spj.science.org/doi/10.1186/s40824-018-0147-5 (accessed on 7 March 2026).

3. Rai, M.; Yadav, A.; Gade, A. Silver Nanoparticles as a New Generation of Antimicrobials. Biotechnol. Adv. 2009, 27, 76–83.

4. Klasyfikacja wyrobów medycznych—Ministerstwo Zdrowia—Portal Gov.pl. Available online: https://www.gov.pl/web/zdrowie/klasyfikacja-wyrobow-medycznych (accessed on 7 March 2026).

5. Kaoutzanis, C.; Winocour, J.; Unger, J.; Gabriel, A.; Maxwell, G.P. The Evolution of Breast Implants. Semin. Plast. Surg. 2019, 33, 217–223.

6. Piechota, M. Badanie USG piersi z implantami—Kompleksowa diagnostyka w Krakowie. In Proceedings of the BodyMove Kraków 2025, Kraków, Poland, 12–14 2025; BodyMove Kraków: Kraków, Poland, 2025; pp. 12–14.

7. Gouthami, S.; Hegde, N.P. Automatic Sentiment Analysis Scalability Prediction for Information Extraction Using SentiStrength Algorithm. In Proceedings of Third International Conference on Advances in Computer Engineering and Communication Systems; Springer: Singapore, 2023; Volume 612, pp. 21–30. ISBN 978-981-19-9228-5.

8. Drope, J.; Liber, A.C.; Cahn, Z.; Stoklosa, M.; Kennedy, R.; Douglas, C.E.; Henson, R.; Drope, J. Who's still smoking? Disparities in adult cigarette smoking prevalence in the United States. CA Cancer J. Clin. 2018, 68, 106–115.

9. U.S. Department of Health and Human Services. The Health Consequences of Smoking: 50 Years of Progress. A Report of the Surgeon General; U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health: Atlanta, GA, USA, 2014.

10. Nortey, E.N.; Agyemang, E.F.; Sakyi-Yeboah, E.; Ampomah, O.-A.; Agyekum, L. AI Meets Economics: Can Deep Learning Surpass Machine Learning and Traditional Statistical Models in Inflation Time Series Forecasting? Data Sci. Financ. Econ. 2025, 5, 136–155.

11. Poghosyan, H.; Moen, E.L.; Kim, D.; Manjourides, J.; Cooley, M.E. Social and Structural Determinants of Smoking Status and Quit Attempts Among Adults Living in 12 US States, 2015. Am. J. Health Promot. 2019, 33, 498–506.

12. Sakko, Y.; Madikenova, M.; Kim, A.; Syssoyev, D.; Mussina, K.; Gusmanov, A.; Zhakhina, G.; Yerdessov, S.; Semenova, Y.; Crape, B.L.; et al. Epidemiology of tuberculosis in Kazakhstan: Data from the Unified National Electronic Healthcare System 2014–2019. BMJ Open 2023, 13, e074208.

13. United States Department of Agriculture Economic Research Service. Food Security in the U.S.: Key Statistics & Graphics. 2021. Available online: https://www.ers.usda.gov/topics/food-nutrition-assistance/food-security-in-the-us/key-statistics-graphics (accessed on 23 March 2025).

14. Pierannunzi, C.; Hu, S.S.; Balluz, L. A systematic review of publications assessing reliability and validity of the Behavioral Risk Factor Surveillance System (BRFSS), 2004–2011. BMC Med. Res. Methodol. 2013, 13, 49. [Central]

15. Mayer, M.; Gueorguieva, R.; Ma, X.; White, M.A. Tobacco use increases risk of food insecurity: An analysis of continuous NHANES data from 1999 to 2014. Prev. Med. 2019, 126, 105765.

16. Snelson, C.L. Qualitative and Mixed Methods Social Media Research: A Review of the Literature. Int. J. Qual. Methods 2016, 15, 1609406915624574.

17. Zhang, T.; Patil, S.G.; Jain, N.; Shen, S.; Zaharia, M.; Stoica, I.; Gonzalez, J.E. RAFT: Adapting Language Model to Domain Specific RAG. arXiv 2024, arXiv:2403.10131.

18. Creswell, J.W.; Poth, C.N. Qualitative Inquiry and Research Design: Choosing Among Five Approaches, 4th ed.; SAGE Publications: Thousand Oaks, CA, USA, 2016; ISBN 978-1-5063-6117-8.

19. Centers for Disease Control and Prevention. Statistical Brief on the Social Determinants of Health and Health Equity Module, Behavioral Risk Factor Surveillance System. 2022. Available online: https://restoredcdc.org/www.cdc.gov/brfss/data_documentation/pdf/SDOH-Module-Statistical-Brief-508c.pdf (accessed on 10 September 2024).

20. Chung-Hall, J.; Fong, G.T.; Meng, G. Evaluating the impact of menthol cigarette bans on cessation and smoking behaviors in Canada: Longitudinal findings from the Canadian arm of the 2016-18 ITC Four Country Smoking and Vaping Surveys. Tob. Control. 2022, 31, 556–563.

21. Zhussupov, B.; Hermosilla, S.; Terlikbayeva, A.; Aifah, A.; Ma, X.; Zhumadilov, Z.; Abildayev, T.; Darisheva, M.; Berikkhanova, K. Risk Factors for Primary Pulmonary TB in Almaty Region, Kazakhstan: A Matched Case-Control Study. Iran. J. Public Health 2016, 45, 441–450.

22. Centers for Disease Control and Prevention (CDC). National Center for Health Statistics (NCHS). In National Health Interview Survey; U.S. Department of Health and Human Services, CDC: Hyattsville, MD, USA, 2023.