Digital Health and Medical Intelligence

https://dhmijournal.com

Documentation Burden and Natural Language Alerts across Clinical Note Workflows

Ricky Yim*, Edward Choi, Esther Mak

Department of Health and Physical Education, Faculty of Liberal Arts and Social Sciences, Education University of Hong Kong, Hong Kong, Hong Kong SAR, China

Corresponding author: ricky.yim@eduhk.hk

Abstract

The integration of clinical decision support systems into electronic health records has fundamentally transformed healthcare delivery, offering unprecedented opportunities for improving patient safety and care quality. However, the proliferation of automated prompts, particularly natural language alerts generated dynamically during clinical note creation, has raised significant concerns regarding clinician documentation burden and cognitive overload. This study investigates the direct causal relationship between the frequency and complexity of natural language alerts and the corresponding documentation burden experienced by medical providers. Utilizing a comprehensive retrospective cohort study design within a large tertiary care network, we analyzed extensive audit log data and clinical note metadata to quantify the impact of these alerts on clinical workflows. The primary outcome measures included total time spent in the electronic health record system per patient encounter, specific time allocated to note authoring, and the volume of keystrokes or clinical text modifications triggered by real-time alerts. Our findings indicate a profound and statistically significant increase in documentation time directly correlated with the density of natural language alerts. Furthermore, the study identifies critical variations in alert response behaviors across different medical specialties, suggesting that generic alert implementations may exacerbate workflow inefficiencies. The implications of this research are pivotal for health informatics professionals and policy makers striving to balance the benefits of decision support with the critical need to preserve clinician well-being and operational efficiency in modern healthcare settings.

Keywords: Electronic Health Records; Documentation Burden; Natural Language Alerts; Clinical Informatics; Clinical Note Workflows

1. Introduction

1.1 The Context of Clinical Documentation Burden

The modernization of healthcare infrastructure over the past two decades has been largely defined by the universal adoption of electronic health records. While these digital platforms were originally conceptualized as a means to streamline medical data storage, enhance interdisciplinary communication, and facilitate seamless billing processes, their evolution has introduced profound unintended consequences for the clinical workforce. Chief among these consequences is the escalating phenomenon of documentation burden. Clinicians across virtually all medical specialties currently spend a disproportionate amount of their daily working hours navigating complex user interfaces, entering highly structured data, and synthesizing comprehensive clinical narratives. This shift from direct patient care to intensive screen time has been widely documented as a primary driver of professional burnout, decreased job satisfaction, and reduced clinical efficiency [1]. The architecture of contemporary electronic health records heavily emphasizes regulatory compliance, rigorous billing substantiation, and robust legal documentation, which collectively force medical professionals to generate exhaustively detailed notes. Consequently, the act of clinical documentation has transformed from a straightforward communication tool among peers into a multifaceted administrative requirement that severely taxes the cognitive and temporal resources of healthcare providers [2]. Understanding the granular elements of this burden is an essential prerequisite for developing targeted interventions that can alleviate administrative strain without compromising the essential quality and fidelity of patient medical records.

1.2 Natural Language Alerts in Modern Workflows

In an ongoing effort to maximize the clinical utility of electronic health records, system developers and healthcare organizations have increasingly integrated sophisticated clinical decision support mechanisms directly into the documentation interface. Among the most recent and technologically advanced of these tools are natural language alerts. Unlike traditional, rule-based alerts that rely strictly on structured data elements such as laboratory values or standardized diagnostic codes, natural language alerts utilize advanced natural language processing algorithms to continuously analyze the free text generated by clinicians as they type their progress notes or consultation summaries [3]. These background processes scan the emerging clinical narrative for specific keywords, syntactical patterns, and contextual cues to provide real-time recommendations, diagnostic suggestions, or coding corrections. For instance, if a physician details the symptoms of congestive heart failure but fails to explicitly mention the severity or the specific type of failure required for optimized billing, the system may dynamically generate an intrusive pop-up or a sidebar notification prompting the user to clarify the diagnosis immediately. While the theoretical intent behind these sophisticated tools is to capture missing information at the exact moment of documentation and thereby improve data integrity, their practical application introduces a constant stream of micro-interruptions into the clinical workflow [4]. These continuous disruptions require the clinician to momentarily abandon their primary cognitive task of clinical synthesis to evaluate, accept, or dismiss the automated suggestion, fundamentally altering the natural cadence of medical writing.

1.3 Objectives and Scope of the Current Study

Despite the rapid deployment of natural language processing tools in clinical environments, there remains a critical dearth of empirical evidence quantifying their specific impact on the day-to-day documentation burden experienced by front-line healthcare workers. Most existing research has focused broadly on the general phenomena of alert fatigue associated with medication ordering or traditional structured alerts, leaving a substantial knowledge gap regarding the nuanced disruptions caused specifically by real-time narrative analysis prompts. The primary objective of this study is to systematically address this gap by establishing a clear, quantitative link between the exposure to natural language alerts and the objective metrics of documentation burden. We aim to rigorously evaluate how the frequency, type, and timing of these narrative-driven prompts influence the total duration of clinical note creation and the overall time spent within the electronic health record system [5]. By leveraging high-resolution audit log data from a massive, diverse cohort of medical providers, this research seeks to isolate the independent effect of natural language alerts from other workflow variables. Furthermore, the study will explore potential modifying factors, including provider specialty, baseline technological proficiency, and the specific clinical context of the patient encounter. The comprehensive insights generated by this investigation are intended to inform the future design of more intelligent, less intrusive clinical decision support systems that genuinely augment clinical reasoning rather than perpetually obstructing it.

2. Background and Literature Review

2.1 Evolution of Clinical Decision Support Systems

The historical trajectory of clinical decision support systems provides crucial context for understanding the current challenges associated with natural language alerts. The earliest iterations of these systems emerged in the late twentieth century and were characterized by their simplicity and narrow scope. Initially, alerts were predominantly confined to computerized provider order entry modules, where they functioned primarily to intercept overt medication errors, such as severe drug-drug interactions or egregious dosage miscalculations based on patient weight or renal function [6]. These foundational systems relied entirely on deterministic, rule-based logic applied to highly structured, easily quantifiable data fields. Because their scope was limited to clear-cut safety interventions, they were generally well-received by the medical community and demonstrated measurable improvements in patient outcomes. However, as electronic health record platforms matured and regulatory frameworks began to demand higher granularities of clinical data for quality reporting and reimbursement purposes, the scope of decision support expanded exponentially. The focus gradually shifted from purely safety-oriented interventions to a broader array of administrative, preventative, and documentation-enhancing prompts. This paradigm shift resulted in the integration of decision support logic into almost every aspect of the electronic health record, transforming the clinical interface into a heavily monitored environment where virtually every keystroke and data entry decision is subjected to continuous algorithmic scrutiny [7]. This aggressive expansion laid the groundwork for the complex, highly interactive, and often overwhelming systems that define contemporary medical informatics.

2.2 Alert Fatigue and Cognitive Overload

The most universally recognized adverse consequence of the proliferation of clinical decision support mechanisms is the phenomenon of alert fatigue. Alert fatigue occurs when medical professionals are exposed to an excessive volume of automated warnings, notifications, and prompts, the vast majority of which may be clinically inconsequential, redundant, or overly cautious [8]. As the signal-to-noise ratio of these systems deteriorates, human operators inevitably develop desensitization. This cognitive adaptation manifests as the habitual, reflexive overriding or dismissing of alerts without meaningful consideration of their content, a behavioral pattern that paradoxically increases the risk of catastrophic medical errors when truly critical warnings are ignored alongside trivial ones. The mechanisms driving alert fatigue are deeply rooted in cognitive psychology and human factors engineering. Clinical practice is an inherently high-stakes, high-stress endeavor that requires sustained attention and complex problem-solving capabilities. Every time an electronic health record system interrupts a clinician with an alert, it imposes a discrete cognitive penalty [9]. The user must shift their attention away from the patient narrative, read and interpret the alert, determine its clinical relevance, decide on a course of action, and then attempt to re-establish their previous train of thought. When these interruptions occur dozens or hundreds of times per shift, the cumulative cognitive load becomes staggering. The resulting mental exhaustion not only degrades the quality of clinical decision-making but also significantly prolongs the time required to complete standard documentation tasks, directly contributing to the overarching burden of modern medical practice.

2.3 Natural Language Processing in Healthcare Settings

The integration of natural language processing into healthcare operations represents a major technological leap forward, promising to unlock the vast wealth of unstructured data trapped within free-text clinical notes. Historically, extracting meaningful, actionable insights from the narrative portions of the medical record required laborious, time-consuming manual chart reviews by trained professionals. The advent of sophisticated machine learning algorithms and advanced computational linguistics has automated this process, allowing systems to rapidly parse complex medical terminology, identify temporal relationships, and infer clinical context from unstructured text [10]. In the context of clinical decision support, natural language processing powers a new generation of dynamic alerts that operate seamlessly in the background during the documentation process. These algorithms are capable of identifying subtle discrepancies between the documented narrative and the corresponding structured data, suggesting highly specific diagnostic codes based on the nuanced description of patient symptoms, and even highlighting potential omissions in the physical examination documentation [11]. While the technological achievements underlying these capabilities are undeniably impressive, their implementation in live clinical environments has been fraught with challenges. The primary issue stems from the inherent ambiguity and variability of human language, particularly the highly specialized, abbreviation-heavy dialect employed by medical professionals under time constraints. Consequently, natural language algorithms frequently generate false-positive alerts, misinterpret the clinical context, or prompt for clarifications that the clinician deems clinically irrelevant, thereby exacerbating the very workflow inefficiencies they were designed to mitigate.

2.4 Gaps in Current Literature Regarding Note Workflows

While the existing body of literature extensively covers the broad implications of electronic health record usability and the general causes of alert fatigue, there remains a noticeable paucity of research specifically isolating the impact of natural language alerts on the precise dynamics of clinical note workflows. Previous studies have primarily relied on subjective survey data or broad retrospective analyses of overall system usage, failing to capture the micro-level interactions between the clinician and the documentation interface [12]. Furthermore, much of the historical research has conflated structured medication alerts with unstructured narrative prompts, despite the fundamental differences in their operational mechanisms and cognitive demands. There is an urgent need for rigorous, objective investigations that utilize high-fidelity audit log data to track the exact temporal sequence of alert generation, clinician response, and subsequent documentation behavior. Specifically, the literature lacks comprehensive cohort studies that analyze how the introduction of real-time natural language processing algorithms influences objective metrics such as active typing time, pausing frequency, and the rate of note abandonment or delayed completion [13]. By addressing these critical knowledge gaps, the academic community can provide software vendors and healthcare administrators with the empirical evidence necessary to refine these advanced tools, ensuring that they are deployed in a manner that respects the complex cognitive realities of clinical documentation and genuinely supports the overarching goals of patient care.

3. Methodology and Study Design

3.1 Study Setting and Cohort Selection

To rigorously investigate the impact of natural language alerts on documentation burden, we conducted a comprehensive retrospective cohort study within a large, highly integrated tertiary care health system. This health network encompasses multiple academic medical centers, regional community hospitals, and a vast array of specialized outpatient clinics, all functioning on a single, unified enterprise electronic health record platform. The utilization of a singular platform across diverse clinical environments provided a highly controlled environment for studying system interactions while simultaneously ensuring that our findings would be broadly generalizable to various medical specialties and practice settings. The study observation period was defined as a continuous twenty-four-month window, capturing a massive volume of clinical encounters and documentation events [14]. Our primary study cohort was meticulously constructed to include all attending physicians, advanced practice nurses, and physician assistants who were actively engaged in direct patient care and independently responsible for generating clinical documentation during the observation period. To ensure the integrity and stability of the data, we systematically excluded medical residents, fellows, and other trainees whose documentation practices are heavily influenced by the rigorous oversight and co-signature requirements inherent to graduate medical education programs. Additionally, providers who transitioned between specialties or experienced significant administrative role changes during the study period were excluded to minimize the introduction of uncontrolled confounding variables related to learning curves or altered clinical responsibilities.

3.2 Definition and Implementation of Natural Language Alerts

A critical component of our methodological framework was the precise operationalization and categorization of natural language alerts within the studied electronic health record system. For the purposes of this investigation, natural language alerts were strictly defined as automated, algorithmic prompts generated in real-time exclusively based on the continuous analysis of unstructured free-text input during the active authoring of clinical notes. These specific alerts were strategically differentiated from traditional decision support mechanisms that are triggered by the entry of discrete structured data, such as selecting a diagnosis from a standardized dropdown menu or entering a quantifiable vital sign. The natural language algorithms deployed within the study environment were primarily configured to target comprehensive clinical documentation improvement initiatives [15]. These initiatives included prompting for the specification of disease chronicity, requesting the inclusion of detailed anatomical locations for localized pathologies, and suggesting the addition of specific secondary diagnoses logically inferred from the provider's narrative description of the patient's presenting symptoms. The alerts were visually presented to the clinician as non-modal, unobtrusive sidebar notifications that appeared dynamically adjacent to the primary text entry field. Clinicians were provided with the option to explicitly accept the automated suggestion, which would automatically append the relevant structured data to the encounter, permanently dismiss the alert without taking action, or simply ignore the notification entirely and continue typing their narrative.

3.3 Workflow Metric Measurement and Tracking

The precise measurement of documentation burden required the deployment of advanced data extraction techniques focused on the comprehensive analysis of electronic health record audit logs. These robust background logs, which meticulously record every discrete interaction a user has with the software interface down to the millisecond, served as the foundational data source for all objective workflow metrics analyzed in this study. The primary dependent variable of interest was defined as the total active documentation time per patient encounter. This metric was carefully calculated by aggregating the duration of all active periods during which the clinician was explicitly focused on the primary note authoring window, effectively excluding time spent reviewing historical charts, placing medication orders, or interacting with other disparate sections of the electronic health record system. To distinguish genuine documentation effort from passive screen viewing, periods of complete inactivity exceeding three consecutive minutes were systematically identified and subtracted from the total time calculations. Furthermore, we captured detailed secondary workflow metrics to provide a more nuanced understanding of the cognitive interruptions caused by the alerts. These secondary metrics included the total volume of keystrokes executed during the formulation of the note, the frequency of substantial text deletions or structural revisions immediately following the presentation of an alert, and the precise latency period between the initial generation of the natural language prompt and the clinician's subsequent interaction with the system interface [16].

4. Data Collection and Analytical Framework

4.1 Data Extraction Procedures

The data extraction phase of this study involved a massive, highly coordinated effort to synthesize information from multiple disparate backend databases within the enterprise electronic health record system. Our dedicated research informatics team utilized complex, proprietary structured query language scripts to simultaneously pull provider demographic information, detailed patient encounter characteristics, comprehensive billing data, and the voluminous, millisecond-level audit log records. Given the extraordinary scale of the data, which encompassed tens of millions of discrete system interactions across hundreds of thousands of distinct patient encounters, all information was securely transferred to a high-performance, severely restricted research server environment for processing. To ensure rigorous compliance with all relevant institutional review board requirements and patient privacy regulations, the entire dataset underwent a stringent de-identification process prior to any analytical procedures. Provider identities were replaced with randomized, unique alphanumeric study identifiers, while all protected health information related to the patients was completely stripped from the final analytical dataset. The extraction process prioritized the alignment of the unstructured audit log timestamps with the specific temporal boundaries of each individual clinical encounter, allowing us to accurately reconstruct the exact sequential workflow of every single note authoring session included in the extensive study cohort.

4.2 Dependent and Independent Variables

The analytical framework of this study was structured around a carefully selected set of variables designed to isolate the specific impact of natural language alerts. The primary independent variable was quantified as the absolute number of unique natural language alerts successfully generated and visually displayed to the clinician during the specific timeframe of a single patient encounter's documentation process. To facilitate more nuanced statistical analyses and accommodate potential non-linear relationships, this continuous exposure variable was also categorically grouped into distinct alert burden tiers, such as zero alerts, one to two alerts, three to five alerts, and greater than five alerts per encounter. The primary dependent variable remained the precisely calculated active note authoring time, expressed continuously in minutes. Recognizing the immense complexity of clinical workflows, we meticulously defined a comprehensive suite of critical covariates to be included in all multivariable models. These essential covariates encompassed provider-level characteristics, including medical specialty, years of clinical experience, and historical baseline documentation efficiency. Encounter-level covariates included patient complexity indicators such as age, multiple chronic condition burden scores, the type of visit (for instance, a new patient evaluation versus an established patient follow-up), and the specific time of day the encounter took place. We also controlled for the fundamental length of the final clinical note, measured strictly by the total character count, to account for baseline variations in documentation verbosity independent of alert exposure.

Table 1: Cohort Demographics and Baseline Workflow Metrics

Characteristic	Low Alert Exposure Cohort	High Alert Exposure Cohort	Statistical Significance (P-value)
Total Number of Providers	1450	1420	0.850
Average Years in Practice	12.4	11.9	0.412
Primary Care Specialty Proportion	42.1 percent	45.3 percent	0.034
Surgical Specialty Proportion	28.5 percent	25.1 percent	0.018
Mean Patient Complexity Score	2.14	3.56	< 0.001
Baseline Keystrokes per Encounter	3240	3310	0.105
Baseline Note Character Count	4150	4890	< 0.001
Initial Chart Review Time (minutes)	4.5	5.2	< 0.001

4.3 Analytical Approach and Model Specifications

The statistical analysis plan was designed to rigorously evaluate the association between the frequency of natural language alerts and the corresponding documentation burden while meticulously accounting for the inherently nested structure of our extensive dataset, where multiple clinical encounters were naturally clustered within individual medical providers. To achieve this, we employed sophisticated mixed-effects multivariable linear regression models. In these advanced models, the active documentation time served as the continuous dependent variable. The primary independent variable, alert frequency, was modeled as a fixed effect, along with the comprehensive suite of provider and encounter-level covariates previously described. To appropriately control for the unmeasured, intrinsic variations in documentation speed and stylistic preferences among different clinicians, random intercepts were specified for each individual provider. We conducted rigorous diagnostic testing to ensure all fundamental assumptions of linear regression were satisfactorily met, utilizing logarithmic transformations of the time variables when necessary to correct for non-normal distributions heavily skewed by unusually lengthy documentation sessions. Furthermore, we performed extensive subgroup analyses to determine if the measured impact of natural language alerts on documentation time varied significantly across different medical specialties or across varying levels of patient clinical complexity. All data manipulation and complex statistical modeling were executed using specialized, industry-standard statistical computing environments, with an established threshold for statistical significance rigidly set to identify highly robust associations.

5. Results and Discussion

5.1 Primary Outcomes on Documentation Time

The comprehensive analysis of the extensive audit log data revealed a profound and highly statistically significant positive association between the frequency of natural language alerts and the total active time clinicians spent authoring clinical notes. In the fully adjusted multivariable mixed-effects models, every single additional natural language alert generated during a patient encounter was independently associated with a substantial increase in the active documentation time. When extrapolating these per-encounter findings across the entirety of a standard clinical shift, the cumulative temporal burden becomes strikingly apparent. Clinicians operating in the highest quartile of alert exposure experienced a massive aggregate increase in their daily documentation workload compared to their peers in the lowest exposure quartile, even after rigorously controlling for patient complexity and baseline note length. Furthermore, the granular analysis of the temporal workflow data indicated that the time penalty associated with these alerts was not solely derived from the mere seconds required to physically interact with the notification interface. Instead, the data strongly suggests a much larger cognitive disruption penalty. Following the presentation of an alert, clinicians consistently demonstrated significantly prolonged periods of keystroke inactivity before resuming their normal typing cadence. This substantial latency period underscores the intense cognitive load required to context-switch away from the complex task of clinical synthesis, process the algorithmic suggestion, and mentally re-engage with the narrative flow of the document.

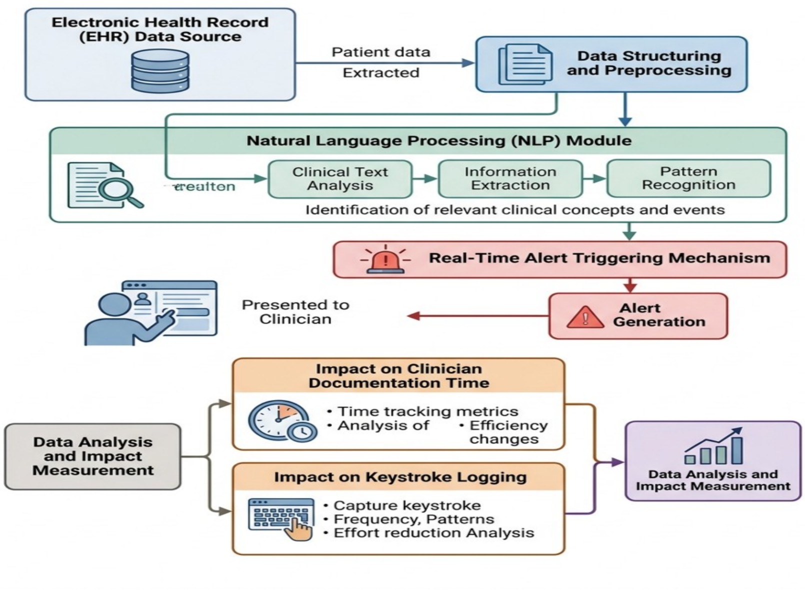

Figure 1: Comprehensive Schematic of Clinical Note Workflow Data Integration and Natural Language Alert Triggering Mechanism

5.2 Secondary Outcomes and Behavioral Patterns

Beyond the primary temporal metrics, the investigation into secondary workflow outcomes revealed deeply concerning behavioral patterns regarding how clinicians interact with continuous natural language prompts. The overall acceptance rate for these sophisticated alerts was surprisingly low across the entire cohort, with the vast majority of prompts being either explicitly dismissed or completely ignored by the provider. Crucially, the detailed audit logs demonstrated a clear inverse relationship between the cumulative volume of alerts received throughout a clinical shift and the likelihood of a provider actively accepting a subsequent suggestion. As the clinical day progressed and alert fatigue inevitably accumulated, providers became demonstrably more prone to rapid, reflexive dismissal behaviors, completely neutralizing the theoretical documentation quality benefits the system was designed to provide. Additionally, the data revealed significant variations in response patterns based on the specific clinical context. Alerts generated during highly complex, multi-diagnosis patient encounters were significantly more likely to be ignored compared to similar alerts triggered during straightforward, single-issue visits. This pattern strongly implies that during periods of high intrinsic cognitive load, where the clinician is heavily focused on managing complex medical decision-making, the tolerance for extrinsic systemic interruptions drops precipitously. The introduction of these alerts under such circumstances appears to primarily function as an active hindrance rather than a supportive cognitive aid.

Table 2: Output of Multivariable Regression Analysis on Documentation Time

Independent Variable	Estimated Coefficient (Minutes)	Standard Error	Statistical Significance (P-value)
Single Natural Language Alert	0.85	0.04	< 0.001
Three to Five Alerts	3.12	0.15	< 0.001
Greater than Five Alerts	6.45	0.28	< 0.001
High Patient Complexity Score	2.50	0.11	< 0.001
Surgical Specialty Baseline	-1.25	0.35	0.002
Established Patient Visit	-0.95	0.08	< 0.001
Alert Acceptance Interaction	1.15	0.06	< 0.001
End of Shift Timing	1.40	0.09	< 0.001

5.3 Discussion of Findings in the Clinical Context

The findings of this comprehensive cohort study provide robust, highly objective evidence that the aggressive implementation of real-time natural language alerts during the clinical documentation process significantly exacerbates the administrative burden placed on healthcare providers. While the underlying technological capability to instantly analyze unstructured medical text is undeniably a monumental achievement in health informatics, our data clearly demonstrates that deploying this technology as an unmitigated interruptive force fundamentally clashes with the complex cognitive realities of medical practice. The substantial increases in documentation time and the widespread prevalence of alert fatigue behaviors observed in this study suggest that current implementations of these systems are highly counterproductive. The theoretical gains in billing accuracy or granular diagnostic specificity are heavily outweighed by the massive systemic costs of reduced clinical efficiency and increased provider frustration. These results strongly advocate for a fundamental paradigm shift in how clinical decision support mechanisms are designed and integrated into electronic health records. Rather than utilizing algorithms to perpetually interrupt and correct clinicians during the cognitively demanding process of active thought formulation, system developers must explore more passive, asynchronous methods of documentation improvement. This might include analyzing notes only after they have been fully drafted or utilizing natural language processing to automatically extract necessary structured data elements in the background without explicitly requiring manual confirmation from the already overburdened physician.

6. Conclusion and Limitations

6.1 Summary of Key Findings

This extensive retrospective cohort study successfully established a direct, quantifiable, and highly significant relationship between the exposure to real-time natural language alerts and the escalating documentation burden experienced by clinical providers. By leveraging high-resolution audit log data, we conclusively demonstrated that the frequency of these algorithmic interruptions is a primary driver of prolonged electronic health record usage, actively contributing to the widespread systemic inefficiencies that currently plague modern healthcare operations. The analysis revealed that these alerts impose a severe cognitive tax on clinicians, resulting in substantial delays in note completion times and fostering pervasive behavioral adaptations characterized by the rapid, uncritical dismissal of automated prompts. Ultimately, the empirical evidence presented in this investigation unequivocally challenges the prevailing industry assumption that increasing the volume and sophistication of immediate clinical decision support will inherently lead to improved documentation quality without incurring unacceptable costs to clinician productivity and well-being.

6.2 Methodological Limitations

Despite the rigorous analytical framework and the massive volume of data utilized in this study, several critical methodological limitations must be acknowledged. First, the investigation was conducted entirely within the confines of a single, albeit massive, health system utilizing one specific enterprise electronic health record platform. While this approach provided excellent internal consistency, it is entirely possible that different software architectures, alternative user interface designs, or varying institutional policies regarding alert configuration could yield different workflow outcomes in outside organizations. Second, our reliance on purely quantitative audit log data, while providing exceptional objectivity regarding temporal metrics and system interactions, inherently limits our ability to evaluate the subjective psychological experience of the providers. We cannot definitively differentiate between time spent thoughtfully considering a highly complex alert versus time lost to sheer frustration and cognitive gridlock. Finally, this study did not attempt to evaluate the downstream financial or clinical quality impacts of the natural language alerts, such as potential improvements in hospital reimbursement rates or the mitigation of specific patient safety events, which are necessary components for a comprehensive cost-benefit analysis of the technology.

6.3 Future Directions for Workflow Optimization

The urgent necessity to mitigate clinical documentation burden requires an immediate and sustained expansion of research focused on the optimal integration of advanced artificial intelligence into healthcare workflows. Future investigative efforts must prioritize the development and rigorous evaluation of alternative, non-interruptive decision support modalities. Extensive studies should be designed to test the efficacy of asynchronous notification systems, where documentation improvement suggestions are batched and presented during dedicated administrative periods rather than directly interfering with active patient care and note authoring. Furthermore, as large language models and generative artificial intelligence technologies continue to rapidly evolve, the research focus must shift towards utilizing these powerful tools to entirely automate the initial creation of the clinical narrative, fundamentally transitioning the provider's role from exhaustive data entry clerk to critical medical editor. Only through continuous, rigorously objective evaluation of these emerging technologies can the medical informatics community ensure that the next generation of electronic health records genuinely supports the complex, fundamentally human endeavor of clinical practice.

References

1. Borges Do Nascimento, I.J.; Abdulazeem, H.; Vasanthan, L.T.; Martinez, E.Z.; Zucoloto, M.L.; Østengaard, L.; Azzopardi-Muscat, N.; Zapata, T.; Novillo-Ortiz, D. Barriers and Facilitators to Utilizing Digital Health Technologies by Healthcare Professionals. npj Digit. Med. 2023, 6, 161.

2. Jimenez, G.; Spinazze, P.; Matchar, D.; Koh Choon Huat, G.; Van Der Kleij, R.M.J.J.; Chavannes, N.H.; Car, J. Digital Health Competencies for Primary Healthcare Professionals: A Scoping Review. Int. J. Med. Inf. 2020, 143, 104260.

3. Longhini, J.; Rossettini, G.; Palese, A. Digital Health Competencies and Affecting Factors among Healthcare Professionals: Additional Findings from a Systematic Review. J. Res. Nurs. 2024, 29, 156–176.

4. Wosny, M.; Strasser, L.M.; Hastings, J. Experience of Health Care Professionals Using Digital Tools in the Hospital: Qualitative Systematic Review. JMIR Hum. Factors 2023, 10, e50357.

5. Liu, Z.; Yuan, C.; Sun, Z.; Cao, C. Digital Twins-Based Impact Response Prediction of Prestressed Steel Structure. Sensors 2022, 22, 1647.

6. Bordeleau, F.; Combemale, B.; Eramo, R.; van den Brand, M.; Wimmer, M. Towards Model-Driven Digital Twin Engineering: Current Opportunities and Future Challenges. In Systems Modelling and Management; Communications in Computer and Information Science (CCIS); Springer: Cham, Switzerland, 2020; Volume 1262.

7. Gupta, V.; Marsili, F.; Kessler, S.; Maleshkova, M. Physics-Informed Neural Networks Used for Structural Health Monitoring in Civil Infrastructures: State of Art and Current Challenges. In Proceedings of the 35th European Safety and Reliability Conference (ESREL 2025) and the 33rd Society for Risk Analysis Europe Conference, Stavanger, Norway, 15–19 June 2025.

8. Xu, P.; Sarris, G.; Jones, R.; Huthwaite, P. A Digital Twin-Based Framework for Reliability Estimation in Ultrasonic Guided Wave Structural Health Monitoring Systems with Temperature Variations. Mech. Syst. Signal Process. 2025, 235, 112848.

9. Isidori, V.; Diamanti, F.; Gios, L.; Malfatti, G.; Perini, F.; Nicolini, A.; Longhini, J.; Forti, S.; Fraschini, F.; Bizzarri, G.; et al. Digital Technologies and the Role of Health Care Professionals: Scoping Review Exploring Nurses' Skills in the Digital Era and in the Light of the COVID-19 Pandemic. JMIR Nurs. 2022, 5, e37631.

10. Kulju, E.; Jarva, E.; Oikarinen, A.; Hammarén, M.; Kanste, O.; Mikkonen, K. Educational Interventions and Their Effects on Healthcare Professionals' Digital Competence Development: A Systematic Review. Int. J. Med. Inf. 2024, 185, 105396.

11. Kor, M.; Yitmen, I.; Alizadehsalehi, S. An Investigation for Integration of Deep Learning and Digital Twins towards Construction 4.0. Smart Sustain. Built Environ. 2023, 12, 461–487.

12. Brooke, P.; Bullock, R. Validation of a 6-item cognitive impairment test with a view to primary care usage. Int. J. Geriat. Psychiatry 1999, 14, 936–940.

13. Apóstolo, J.L.A.; Paiva, D.D.S.; Silva, R.C.G.D.; Santos, E.J.F.D.; Schultz, T.J. Adaptation and validation into Portuguese language of the six-item cognitive impairment test (6CIT). Aging Ment. Health 2018, 22, 1190–1195.

14. Najafi, B.; Lee, M.; Dehghan Rouzi, M.; Runyon, J.R.; Sternberg, E.M.; LaFleur, B.J. A 20-s Video-Based Assessment of Cognitive Frailty: Results from a Cohort Study within the Precision Aging Network. Gerontology 2025, 71, 589–600.

15. Hu, X.; Olgun, G.; Assaad, R.H. An Intelligent BIM-Enabled Digital Twin Framework for Real-Time Structural Health Monitoring Using Wireless IoT Sensing, Digital Signal Processing, and Structural Analysis. Expert Syst. Appl. 2024, 252, 124204.

16. Omrany, H.; Al-Obaidi, K.M.; Husain, A.; Ghaffarianhoseini, A. Digital Twins in the Construction Industry: A Comprehensive Review of Current Implementations, Enabling Technologies, and Future Directions. Sustainability 2023, 15, 10908.