Journal of Frontiers in Multidisciplinary Research

AI-Assisted Nursing Handoff Intelligence for Improving Information Continuity Escalation Reliability and Closure of Unresolved Patient-Care Concerns

Taiwo Juliana Dada [1*], Abutu Ann Oine [2]
[1] Department of Nursing Science, Grand Canyon University, Phoenix Arizona, USA
[2] Department of Medicine and Surgery, Belgorod State National Research University, Belgorod, Russia

* Corresponding Author: **Taiwo Juliana Dada**

Article Info

E-ISSN: 3050-9726
P-ISSN: 3050-9718
Impact Factor (RSIF): 8.10
Volume: 05
Issue: 02
July - December 2024
Received: 09-06-2024
Accepted: 07-07-2024
Published: 05-08-2024
Page No: 170-188

Abstract

Incomplete transfer of clinical information during nursing handoffs can allow unresolved patient-care concerns to persist across shifts, delay escalation of deteriorating conditions, and create uncertainty about whether identified risks have been appropriately addressed. This study proposes HandoffCARE-Net, a Handoff Continuity, Alerting, Resolution, and Escalation Network, as an artificial intelligence framework for automatically analysing sequential nursing handoff records and identifying clinically important information that remains unresolved. HandoffCARE-Net combines a clinical transformer-based text encoder, cross-handoff temporal attention, patient-concern state memory, risk-sensitive escalation prioritisation, and a closure-verification layer within a unified multi-task learning architecture. The model simultaneously performs unresolved-concern detection, information-continuity classification, escalation-priority prediction, concern-state tracking, and closure-status verification across consecutive handoffs. A temporal concern graph links symptom, interventions, pending investigations, medication-related issues, deterioration indicators, nursing actions, escalation events, and documented resolutions, enabling the algorithm to distinguish newly identified concerns from continuing, escalated, resolved, and potentially lost concerns. The proposed system is evaluated against conventional structured handoff and checklist-based approaches and computational benchmarks including TF-IDF Support Vector Machine, XGBoost, BiLSTM-Attention, and ClinicalBERT. Comparative performance is assessed using sensitivity, specificity, precision, macro-F1 score, AUROC, area under the precision-recall curve, false-negative rate, calibration error, escalation-detection accuracy, concern-closure accuracy, and processing latency. Additional ablation experiments examine the contribution of temporal memory, escalation weighting, concern-graph modelling, and closure verification. Performance graphs include algorithm-level ROC and precision-recall curves, sensitivity and F1 comparisons, calibration plots, unresolved-concern detection rates, escalation reliability analysis, and component-ablation charts. The central hypothesis is that integrating longitudinal handoff context with explicit escalation and closure reasoning will provide more reliable detection of clinically significant unresolved concerns than isolated text classification or conventional structured handoff methods. HandoffCARE-Net therefore provides a technical foundation for AI-assisted nursing handoff surveillance capable of strengthening information continuity, supporting timely escalation, and improving verification that patient-care concerns are followed through to documented resolution.

DOI: https://doi.org/10.54660/.JFMR.2024.5.1.349-367

Keywords: AI-Assisted Nursing Handoff, Information Continuity, Escalation Reliability, Unresolved Patient-Care Concerns, Handoff Intelligence

1. Introduction

1.1. Background to Nursing Handoff and Clinical Information Continuity

Nursing handoff is the controlled transfer of patient information, clinical responsibility, and anticipated actions from an outgoing nurse to an incoming nurse at a change of shift or care setting. Its safety function is not simply to summarize the patient, but to preserve a usable state across time. Information continuity therefore requires that active diagnoses, physiological trends, recent interventions, pending investigations, medication issues, escalation thresholds, and incomplete nursing tasks remain traceable after responsibility changes. Lockwood (2016) [28] emphasizes that continuity depends on how reliably information survives handover, while Giske *et al*. (2018) [14] show that handoff quality depends on structure, communication behavior, and accurate relay of key information. In systems terms, this resembles a continuity-control problem in which state information must remain available and actionable after ownership changes. Kwarteng *et al*. (2021) [25], although working in distributed organizational systems, similarly demonstrate the importance of coordinated information flow for reliable service delivery. For this study, nursing handoff intelligence is modeled as longitudinal rather than document-level classification. HandoffCARE-Net represents each patient concern as a temporally evolving state that may be newly identified, continuing, escalated, resolved, or lost across successive handoffs. This design addresses a weakness of isolated note analysis: a statement such as "potassium replacement commenced; repeat level pending" cannot be judged safely without determining whether the next nurse received the pending laboratory requirement and whether later documentation confirms completion. Secure continuity is equally important because handoff information may include data from bedside devices and electronic records; Idika (2022) [15] demonstrates the broader relevance of lightweight authentication in protecting medical-device communications. Accordingly, the framework combines contextual clinical-text encoding, cross-handoff temporal attention, concern-state memory, and closure verification. These components establish the conceptual basis for comparisons using information-continuity accuracy, unresolved-concern sensitivity, escalation reliability, false-negative rate, and closure-status classification.

1.2. Patient-Safety Risks Associated with Incomplete Handoff Information

Incomplete handoff information becomes a serious patient-safety hazard when omitted, distorted, delayed, or poorly prioritized details alter the receiving nurse's assessment or action. The risk is important for variables such as abnormal vital signs, deteriorating neurological status, medication changes, fluid balance, pending cultures, laboratory abnormalities, and time-dependent interventions. Mardis *et al*. (2017) [30] found that shift-to-shift handoff interventions have relevance to safety outcomes, while Cross *et al*. (2019) [10] identified vital-sign communication as important because deficient physiological information can impair decisions. From a systems perspective, a handoff omission functions like an observability failure: the underlying patient risk persists, but the receiving clinician may no longer have sufficient information to recognize it. Amebleh and Omachi (2022) [2], in a high-throughput monitoring context, illustrate the transferable principle that early anomaly detection depends on explicit observability, service thresholds, and rapid identification of departures from expected states.

The clinical consequence is not limited to outright omission. Information can be present but functionally incomplete when urgency, ownership, timing, or required follow-up is unclear. For example, "blood pressure low, doctor informed" does not indicate whether review occurred, treatment was initiated, or further escalation remains necessary. HandoffCARE-Net therefore, treats safety-critical information as a linked sequence of concern, action, escalation, and closure rather than as independent sentences. The model assigns risk-sensitive weights to unresolved concerns and evaluates whether essential elements survive into subsequent handoffs. Cybersecurity principles also matter because loss of integrity or inappropriate access can compromise clinical information; Kwarteng *et al*. (2020) [26] show, in organizational IT operations, how access control and security awareness support risk reduction. These ideas motivate metrics emphasizing clinically consequential errors, particularly sensitivity, false-negative rate, escalation-detection accuracy, calibration, and concern-closure accuracy. Such measures permit later comparison of HandoffCARE-Net with checklist, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT baselines.

1.3. Persistence of Unresolved Patient-Care Concerns Across Nursing Shifts

Unresolved patient-care concerns are problems that have been recognized but have not yet reached a credible documented endpoint when responsibility passes to another nurse. They include pending diagnostic results, incomplete medication adjustments, abnormal observations awaiting reassessment, unreviewed deterioration, unresolved pain and interventions not yet verified. Ernst *et al*. (2018) [12] describe nurse-to-nurse handoff as a process embedded within the broader flow of nursing work, showing that shared understanding of patient condition is constructed across preparation, communication, and subsequent care. Thomson *et al*. (2018) [41] show that handoff quality is influenced by workflow, safety climate, interpersonal relationships, and operational conditions. This matter because an unresolved concern can be transmitted verbally yet still become functionally lost if its status, ownership, or next action is not maintained. Atalor (2022) [6], although focused on pharmacovigilance infrastructure, provides a transferable traceability principle: longitudinal safety events require persistent, auditable records across organizational boundaries. This study formalizes unresolved concerns as temporal clinical entities rather than isolated text labels. For a concern cic_i, HandoffCARE-Net estimates its state across consecutive handoffs and determines whether evidence supports continuation, escalation, resolution, or disappearance without closure. A concern such as "urine output falling; review requested" should remain active until subsequent documentation records assessment, intervention, normalization, or another defensible endpoint. Patient-concern memory and temporal graphs preserve this lineage. Cross-domain work on multilingual information transfer shows that communication systems must account for differences in representation and interpretation; Ijiga *et al*. (2021a) [17], though situated in education, illustrates the importance of information structures that remain interpretable across heterogeneous users. The framework directly measures concern-retention rate, unresolved-concern recall, closure accuracy, and temporal consistency. These outcomes are essential because a model may achieve strong document-level classification while still failing to preserve high-risk concerns over multiple nursing shifts.

1.4. Escalation Failure and Delayed Clinical Response

Escalation reliability refers to the probability that a clinically significant concern is recognized, communicated, prioritized by risk, and followed until a response is documented. Failure can occur when deterioration may be missed, concern may be communicated without sufficient urgency, the receiving clinician may delay action, or the event may be handed over without confirmation that escalation produced a response. White *et al.* (2019) [46] showed that nurses can identify deterioration yet still delay escalation to rapid-response or cardiac-arrest teams, demonstrating that recognition and escalation are distinct processes. Ede *et al.* (2020) [11] identified modifiable barriers and facilitators to escalation in acute wards. These findings require more than abnormal-language detection. Atalor (2019) [5], in privacy-preserving adverse-event prediction, illustrates the value of distributed predictive systems that identify safety signals without centralizing sensitive information, a principle relevant to multi-unit clinical deployment. HandoffCARE-Net therefore includes a risk-sensitive escalation module combining semantic concern representations with temporal change, severity cues, prior escalation evidence, and unresolved duration. Its output is not intended to replace nursing judgment; it ranks concerns requiring verification and highlights cases where the documentation chain suggests escalation has stalled. A patient whose respiratory rate rises across two shifts, for example, should receive a higher escalation score when the record contains repeated concern statements but no documented medical review or intervention. Deep-learning security research in distributed architectures, such as Idika *et al.* (2021) [16], also demonstrates how pattern classification can operate across complex event streams, although its malware domain is distinct from clinical care. Analysis will distinguish concern detection from escalation reliability by reporting sensitivity for high-risk concerns, escalation-priority accuracy, false-negative reduction, AUROC, calibration error, and processing latency. This separation is necessary because an algorithm that detects a concern but ranks it below action thresholds may still fail clinically.

1.5. Limitations of Conventional Structured Handoff and Checklist Systems

Structured nursing handoff tools such as SBAR, ISBAR, and checklist-based protocols improve consistency by specifying categories that should be communicated, but standardization does not guarantee that important concerns remain visible after handoff. Müller *et al.* (2018) [31] found evidence that SBAR can improve communication and patient-safety outcomes, while noting study heterogeneity. Malfait *et al.* (2018) [29] demonstrated high compliance with an ISBARR-based bedside handover protocol but also identified omitted elements and contextual influences on execution. This exposes a longitudinal surveillance limitation: a structured form can confirm that a category was addressed without determining whether an unresolved concern was subsequently acted upon. A checklist item stating “pending investigations” may be completed even when the next shift fails to document the result, response, or closure. Ijiga *et al.* (2021b) [18], although concerned with digital storytelling in education, provides a transferable communication insight that structured presentation alone does not ensure equivalent retention across recipients.

The proposed study therefore treats structured handoff and checklist methods as clinical baselines rather than sufficient endpoints. HandoffCARE-Net adds semantic interpretation, cross-shift temporal linkage, escalation scoring, and closure verification to determine whether critical information persists beyond transfer. This distinction resembles the broader difference between formal rule declaration and realized operational effect. Ajayi *et al.* (2019) [1], writing in a legal context, examine whether a formally declared norm becomes practically achievable; the analogy is limited but useful for distinguishing protocol presence from effectiveness. Technically, the study compares structured handoff against TF-IDF SVM, XGBoost, BiLSTM-Attention, ClinicalBERT, and HandoffCARE-Net using identical patient-level splits and safety-oriented outcomes. Graphical analysis includes sensitivity and F1 comparisons, ROC and precision-recall curves, calibration plots, false-negative distributions, and closure-verification performance. The discussion will therefore examine whether longitudinal AI reasoning adds measurable protection against unresolved concerns that structured documentation alone cannot reliably track.

1.6. Artificial Intelligence for Clinical Communication Intelligence

Artificial intelligence can extend clinical communication analysis beyond keyword matching by learning semantic relationships among symptoms, interventions, temporal expressions, pending tasks, and escalation statements embedded in free-text nursing documentation. Koleck *et al.* (2019) [23] showed that natural language processing can systematically extract symptom information from electronic health-record narratives, demonstrating the clinical value of converting unstructured text into analyzable representations. van Buchem *et al.* (2021) [44] similarly reviewed digital-scribe systems using speech recognition and NLP for clinical documentation and consistently found that most work concentrated on technical tasks such as entity extraction, classification, and summarization, with comparatively limited clinical validation. These observations support models evaluated not only on text-processing accuracy but also on clinically meaningful continuity outcomes. Anokwuru *et al.* (2022) [4] further frame human-AI collaboration as cognitive augmentation rather than replacement, a principle directly applicable to nursing handoff intelligence.

HandoffCARE-Net operationalizes this approach through a clinical transformer encoder, cross-handoff temporal attention, patient-concern memory, temporal concern graphs, escalation prioritization, and closure verification. The architecture is designed to recognize semantically equivalent concern descriptions even when wording changes between nurses. For example, “increasing oxygen need,” “O2 requirement now 4 L,” and “desaturating unless on four litres” should map to a continuing respiratory-risk state rather than unrelated text events. The model explicitly preserves human oversight by presenting prioritized concern trajectories for nurse verification rather than autonomously issuing clinical orders. Evidence from AI-enabled e-learning in low-bandwidth settings (Ijiga *et al.*, 2022) [19], although outside healthcare, reinforces the engineering requirement that intelligent systems remain effective under resource constraints. Accordingly, the study includes processing latency, calibration, error analysis, and ablation testing in addition to AUROC and F1. These analyses will show whether gains arise from temporal memory, escalation weighting, concern-graph construction, and closure verification, connecting algorithmic design directly to patient-safety performance.

1.7. Research Problem and Knowledge Gap

The research gap is that existing handoff improvement and clinical NLP approaches address different parts of the safety problem without jointly modelling the lifecycle of a patient-care concern. Handoff tools structure what should be communicated, while many NLP systems extract entities, classify notes, or normalize concepts from individual documents. Kreimeyer *et al*. (2017) [24] identified temporal information extraction and concept normalization among challenges in clinical NLP, and Wang *et al*. (2018) [45] showed that clinical information-extraction applications remain heterogeneous in tasks, datasets, and evaluation strategies. Neither capability alone determines whether a concern identified on one shift remains visible, is appropriately escalated, and reaches documented closure on a later shift. Atalor *et al*. (2023) [7], although focused on computational drug screening, illustrates the broader value of computational models for complex biomedical inference; however, handoff safety requires a temporal workflow-aware formulation rather than generic prediction.

HandoffCARE-Net is proposed to close this gap by representing each concern as a stateful object embedded in a sequence of nursing handoffs. The framework integrates semantic encoding with temporal attention, concern-state memory, graph-based linkage of observations and actions, escalation-risk prediction, and explicit closure verification. This architecture addresses four needs: continuity detection, unresolved-concern identification, escalation reliability, and resolution tracking. Multi-omics work by Imoh and Idoko (2022) [20], though focused on computational biology, demonstrates a transferable principle of integrating heterogeneous signals rather than relying on a single data stream. Here, heterogeneous signals include clinical text, temporal markers, concern categories, action statements, escalation evidence, and documented outcomes. Comparative evaluation against structured handoff, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT will determine whether this integration improves sensitivity, macro-F1, AUROC, calibration, false-negative control, escalation accuracy, and closure accuracy. Component-level ablation will test whether any observed advantage is attributable to the proposed mechanisms rather than model size alone.

1.8. Research Objectives and Research Questions

- **The research objectives are:**

1. to develop HandoffCARE-Net as an AI-assisted longitudinal nursing handoff intelligence algorithm capable of detecting and tracking unresolved patient-care concerns across consecutive shifts.
2. to formulate temporal concern-state modelling mechanisms for distinguishing newly identified, continuing, escalated, resolved, and unclosed patient-care concerns.
3. to develop risk-sensitive mechanisms for predicting escalation priority and verifying documented concern closure.
4. to compare HandoffCARE-Net with structured handoff/checklist methods, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT under identical patient-level experimental conditions.
5. to evaluate the algorithms using sensitivity, specificity, precision, macro-F1, AUROC, AUPRC, false-negative rate, escalation accuracy, closure accuracy, calibration error, and processing latency.
6. to quantify the individual contribution of temporal memory, concern-graph modelling, escalation weighting, and closure verification through ablation analysis.

- **The corresponding research questions are:**

1. How accurately can HandoffCARE-Net detect unresolved patient-care concerns across successive nursing handoffs?
2. Does longitudinal concern-state modelling improve information continuity compared with isolated handoff classification?
3. Can HandoffCARE-Net improve the reliability of identifying concerns requiring escalation and verifying their eventual closure?
4. How does HandoffCARE-Net perform relative to structured handoff systems and established machine-learning and clinical NLP algorithms?
5. Which HandoffCARE-Net components contribute most strongly to predictive performance and false-negative reduction?
6. Can the proposed system achieve improved safety-oriented performance while maintaining acceptable calibration, interpretability, and computational latency for clinical deployment?

1.9. Contributions and Significance of the Study

This study contributes a longitudinal formulation of nursing handoff intelligence in which the principal analytical unit is not an individual sentence or handoff document but the evolving lifecycle of a patient-care concern. HandoffCARE-Net introduces an integrated architecture combining clinical transformer representations, cross-handoff temporal attention, persistent patient-concern memory, temporal concern graphs, risk-sensitive escalation prediction, and explicit closure verification. Its methodological contribution lies in linking concern recognition to what subsequently happens to that concern, thereby distinguishing simple mention detection from clinically meaningful continuity. The study also introduces a multi-dimensional evaluation framework in which conventional discrimination metrics are supplemented by concern-retention rate, high-risk false-negative rate, escalation-priority accuracy, closure accuracy, calibration, latency, and component-level ablation. The comparative design establishes whether improvements arise from genuine longitudinal reasoning rather than greater model complexity by benchmarking HandoffCARE-Net against structured handoff/checklist procedures, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT.

Clinically, the proposed system is significant because it focuses attention on concerns that may disappear between shifts despite remaining unresolved. Rather than replacing nurses or automatically determining treatment, HandoffCARE-Net is designed as a decision-support layer that presents traceable concern trajectories and prioritized safety signals for professional verification. This approach may strengthen accountability across transitions, reduce information attrition, improve escalation reliability, and provide measurable evidence that important concerns have progressed from recognition to appropriate action and documented closure.

1.10. Scope of the Review and Structure of the Paper

The paper focuses on AI-assisted analysis of sequential nursing handoff information in inpatient clinical environments, particularly the preservation, escalation, and resolution of patient-care concerns during shift-to-shift transfer of responsibility. Its technical scope includes free-text nursing narratives, structured handoff elements, temporal markers, clinical concern categories, documented interventions, escalation events, pending actions, and evidence of resolution. HandoffCARE-Net is evaluated as a decision-support and surveillance architecture rather than an autonomous diagnostic or treatment system; consequently, the study does not propose independent prescribing, diagnosis, replacement of nursing judgment, or unsupervised clinical escalation. The literature review in Section 2 examines nursing handoff communication and patient safety, structured handoff and escalation practices, clinical NLP and machine learning, temporal clinical modelling, and the research gap motivating HandoffCARE-Net. Section 3 presents the formal system model, handoff sequence representation, proposed architecture, temporal concern tracking mechanisms, multi-task learning formulation, comparator algorithms, experimental protocol, and evaluation metrics. Section 4 presents comparative graphical and numerical results covering unresolved-concern detection, information continuity, escalation reliability, closure verification, discrimination, calibration, computational performance, and ablation analysis. Section 5 synthesizes the findings, establishes their clinical and technical implications, identifies limitations, and provides recommendations for implementation, electronic health-record integration, human oversight, AI governance, and subsequent multi-institutional validation.

2. Literature Review

2.1. Nursing Handoff Communication, Information Continuity, and Patient Safety

Nursing handoff is a safety-critical transition in which clinical responsibility, patient state, pending actions, and escalation expectations must remain intelligible to the receiving nurse. Communication quality therefore depends not only on completeness at transfer but also on information continuity across subsequent shifts as shown in figure 1. Kim *et al*. (2021) [22] linked handoff quality with patient-safety culture and emphasized standardized guidance, education, and reliable information exchange. Starmer *et al*. (2017) [40] similarly showed that a structured nursing handoff bundle can improve communication quality without disrupting workflow. For AI-assisted handoff intelligence, these findings imply that a system should preserve the identity and status of important concerns rather than merely summarize the latest note. FHIR-based interoperability principles are relevant because continuity depends on consistent movement of structured and unstructured clinical data across systems (Nwokocha *et al*., 2021a) [32]. Wearable biosensor integration illustrates how physiological evidence can enrich longitudinal monitoring when observations remain connected to decisions (Atalor *et al*., 2023) [8].

Within HandoffCARE-Net, continuity is represented as a temporal sequence linking symptoms, observations, interventions, pending investigations, escalation events, and documented resolution. A concern such as “new confusion with falling oxygen saturation” should remain active until later handoffs provide credible evidence of reassessment, treatment, escalation, or closure. This design treats information loss as a patient-safety failure mode rather than a documentation defect. Healthcare compliance automation also supports traceable, auditable event chains when decisions depend on evidence that required actions occurred (Frimpong *et al*., 2023) [13]. Consequently, the proposed study evaluates continuity using unresolved-concern sensitivity, concern-retention accuracy, escalation-detection reliability, closure accuracy, false-negative rate, and calibration. These measures support comparison of HandoffCARE-Net with structured handoff, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT, where superior performance must reflect preservation of actionable information, not merely higher text-classification accuracy.

Fig 1: Multidisciplinary Bedside Nursing Handoff Supporting Information Continuity, Clinical Coordination, and Patient Safety (Toni, *et al*., 2016).

Figure 1 depicts a multidisciplinary bedside handoff in which several healthcare professionals collectively review a pediatric patient's condition while consulting written clinical records and exchanging observations. This setting illustrates the core functions of nursing handoff communication, information continuity, and patient safety because responsibility for care is being transferred or coordinated through direct verbal discussion supported by documented patient data. The clinicians appear to be comparing charted information, bedside findings, and current care requirements, while another team member records additional details, creating a closed-loop communication process in which information can be confirmed, clarified, and updated immediately. From an information-continuity perspective, this type of handoff should preserve critical elements such as current diagnosis, vital-sign trends, medication status, recent interventions, pending investigations, response to treatment, escalation requirements, and unresolved concerns. Bedside participation also allows the team to validate written information against the patient's present clinical state, reducing the risk that outdated or incomplete documentation is accepted without verification. Technically, an effective handoff in this context depends on accurate transfer of semantic content, temporal sequence, responsibility, and action status: for example, whether an abnormal observation is newly identified, persistent, already escalated, under treatment, or resolved. When these elements are communicated clearly and documented consistently, the receiving team is better positioned to recognize deterioration, avoid duplication or omission of care, maintain continuity across shifts, and ensure that unresolved patient-care concerns remain visible until appropriate clinical closure is achieved.

2.2. Structured Handoff, Checklist-Based Approaches, and Escalation Reliability

Structured handoff frameworks and checklists reduce communication variability by specifying minimum information expected during transfer, including patient situation, background, assessment, recommendations, pending tasks, and anticipated risks. Bukoh and Siah (2020) [9] found that structured nursing handovers can reduce omissions, inaccurate information, and documentation errors, although effects vary across settings as shown in figure 2. Uhm *et al*. (2018) [43] likewise reported improved handover accuracy and fewer handover-related errors after implementing a standardized interdepartmental protocol. These findings establish structured handoff as an appropriate baseline for HandoffCARE-Net, but they also expose an important limitation: checklist completion does not prove that a concern remains visible after transfer or that recommended escalation occurs. Agile system-integration principles emphasize coordinated workflows and consistent information transfer between components, which is relevant when handoff data move across clinical documentation environments (Nwokocha *et al*., 2021b) [33]. Secure microservice deployment research similarly highlights reliable state exchange across distributed components (Ononiwu *et al*., 2023) [34].

Escalation reliability extends beyond communicating that a patient is deteriorating; it requires recognition of severity, assignment of urgency, transfer of responsibility, and verification that a response followed. In HandoffCARE-Net, these steps are modeled through a risk-sensitive escalation score combining concern semantics, temporal persistence, deterioration cues, prior escalation evidence, and closure status. For example, repeated documentation of hypotension without evidence of medical review should increase the priority of that concern across successive shifts. Sustainable healthcare technology design also supports evaluating systems through operational effectiveness rather than technical novelty alone (Anokwuru & Okoh, 2023) [3]. Comparative experiments therefore separate communication completeness from escalation performance using sensitivity, macro-F1, AUROC, false-negative rate, escalation-priority accuracy, and closure verification. This distinction is essential because a structured handoff may capture a risk correctly while still failing to ensure that the risk receives follow-through.

Figure 2 presents structured nursing handoff and escalation reliability as two tightly connected safety mechanisms. The first branch shows how standardized handoff frameworks such as SBAR, ISBAR, I-PASS, bedside protocols, and structured shift reports organize critical information into predictable categories including current condition, diagnosis, vital-sign trends, medications, treatment changes, pending investigations, outstanding tasks, and anticipated risks. Checklist-based approaches improve consistency and reduce omission, but they remain limited because completion of a communication item does not prove that a concern has been clinically resolved. The second branch therefore extends the process into escalation reliability, beginning with recognition of deterioration through abnormal observations, worsening symptoms, laboratory abnormalities, altered consciousness, or increasing treatment requirements. These signals must then pass through severity assessment, prioritization, notification, clinical review, intervention, reassessment, and documented resolution. Failure can occur at any point through delayed recognition, incomplete transfer, unclear responsibility, or absent follow-up. HandoffCARE-Net strengthens this pathway by adding temporal concern tracking, persistent unresolved-concern memory, escalation-priority prediction, and closure verification, thereby converting a static handoff event into a longitudinal safety process in which clinically significant concerns remain visible until appropriate action and documented resolution are achieved.

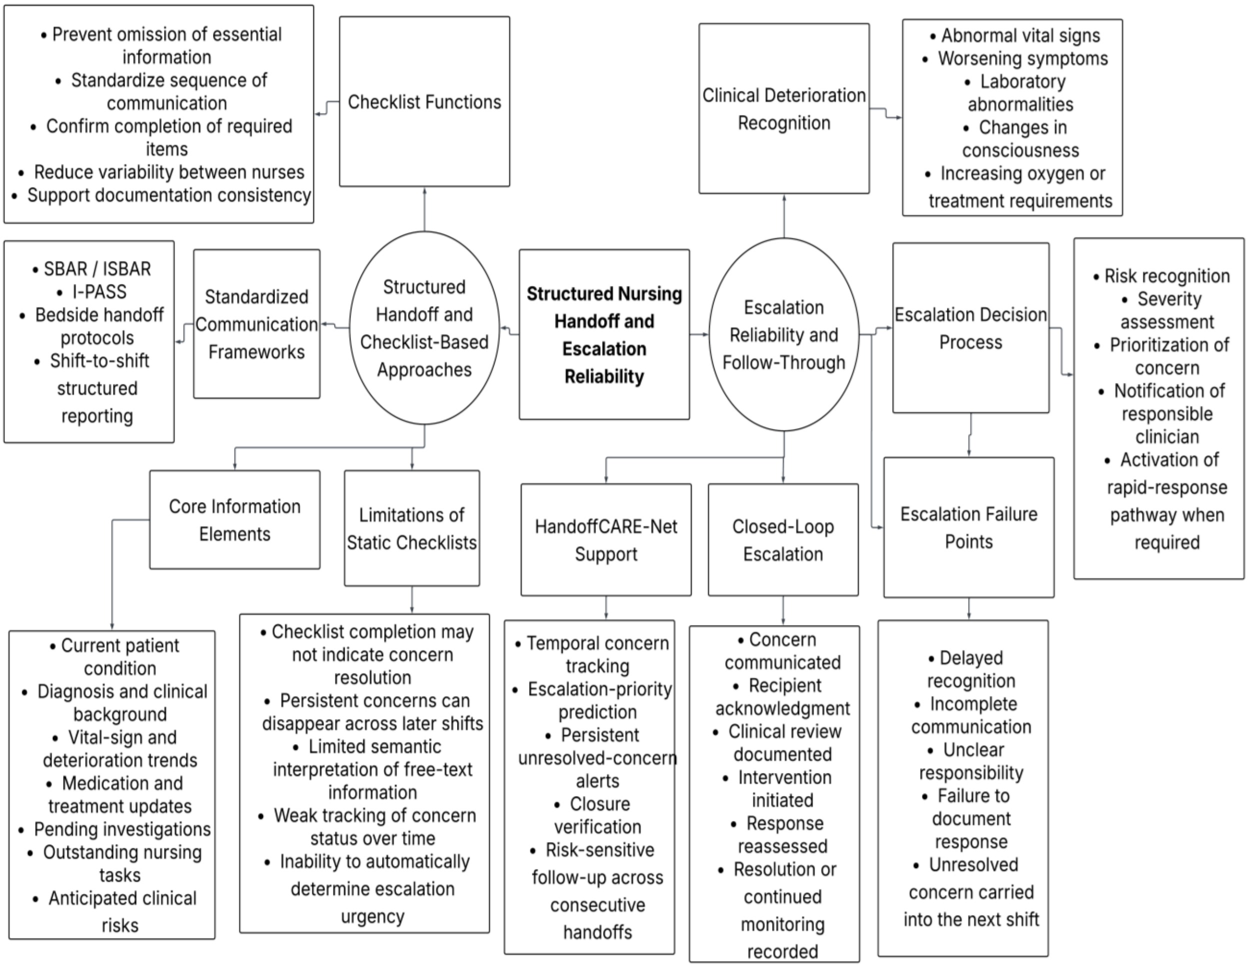

Fig 2: Structured Nursing Handoff, Checklist-Based Communication, and Escalation Reliability Pathway.

2.3. Natural Language Processing and Machine Learning for Clinical Handoff Analysis

Natural language processing and machine learning provide the computational basis for transforming free-text nursing handoffs into structured representations of patient concerns, actions, temporal relationships, and unresolved risks. Wu *et al*. (2020) [48] showed that deep learning is central to clinical NLP, particularly for text classification, named-entity recognition, and relation extraction from electronic health records as shown in figure 3. Percha *et al*. (2022) [38] demonstrated that pretrained language models can infer relevant attributes from unstructured notes, while also revealing errors caused by historical-context confusion, abbreviations, and subtle terminology differences. These limitations are directly relevant to nursing handoffs, where concern may be expressed differently across shifts. AI-enabled medication monitoring provides another example of how machine learning can support prediction and adverse-event recognition from heterogeneous patient information (Onyekaonwu *et al*., 2019) [37]. HandoffCARE-Net therefore uses contextual clinical embeddings rather than isolated keywords to distinguish current, historical, negated, improving, and unresolved concern statements. The proposed architecture combines a transformer encoder with cross-handoff temporal attention, patient-concern memory, temporal concern graphs, escalation prioritization, and closure verification. This design goes beyond single-note classification by estimating how each concern changes state across the handoff sequence. Human-centered intelligent-system design is relevant because decision support must remain interpretable, usable, and subordinate to professional judgment (Onwuzurike, 2023) [35]. Evidence from AI-supported learning systems shows that model performance must be assessed within operational constraints rather than accuracy alone (Onwuzurike & Kpogli, 2022) [36]. Accordingly, HandoffCARE-Net is compared with TF-IDF SVM, XGBoost, BiLSTM-Attention, ClinicalBERT, and structured handoff using precision, sensitivity, specificity, macro-F1, AUROC, AUPRC, calibration error, false-negative rate, and latency. Ablation experiments remove temporal memory, concern-graph linkage, escalation weighting, and closure verification individually, allowing the results to determine whether superior performance derives from the longitudinal mechanisms rather than increased model complexity.

Fig 3: AI-Assisted Natural Language Processing and Machine Learning for Real-Time Nursing Handoff Analysis and Clinical Concern Tracking

Figure 3 illustrates a practical implementation of natural language processing and machine learning within a bedside nursing handoff workflow. The outgoing nurse transfers structured and unstructured clinical information to the incoming nurse while the patient remains physically present, creating a realistic context in which HandoffCARE-Net can process handoff narratives together with vital signs, medication information, pending investigations, and unresolved concerns. The AI handoff interface demonstrates how NLP converts free-text statements into clinically meaningful representations by identifying entities such as shortness of breath, low potassium, pending chest imaging, medications, and physiological measurements, while temporal language processing determines whether these findings are new, persistent, improving, or unresolved across shifts. Machine-learning components then classify concern status, estimate escalation priority, generate a clinical risk score, and assess closure evidence. In the example shown, "shortness of breath" is identified as a high-priority concern persisting across two shifts, potassium of 3.2 is recognized as an unresolved abnormal laboratory finding awaiting replacement, and the chest radiograph remains incomplete. The interface also displays information-continuity status, concern tracking across night, day, and evening shifts, closure verification, and escalation recommendations. Technically, this reflects the proposed HandoffCARE-Net architecture in which contextual transformer embeddings, temporal attention, persistent concern-state memory, and risk-sensitive prediction operate together to transform heterogeneous handoff information into traceable clinical concern trajectories. The human nurses remain central to decision-making, while the AI functions as a surveillance and decision-support layer for reducing information loss, missed escalation, and premature closure of patient-care concerns.

2.4. Transformer Models, Temporal Learning, and Clinical Concern Tracking

Transformer architectures provide a strong basis for nursing handoff intelligence because self-attention can model dependencies among clinically related events separated by multiple observations or care transitions. Li *et al*. (2020) [27] introduced BEHRT as a transformer-based representation of longitudinal electronic health records and demonstrated that attention-based sequence modeling can incorporate heterogeneous clinical concepts while capturing long-range relationships across patient histories. Rasmy *et al*. (2021) [39] extended this direction through Med-BERT, demonstrating that large-scale pretraining on structured health records can generate contextual representations that improve downstream clinical prediction when labeled data are limited. These properties are important for nursing handoffs because concern meaning is often distributed temporally. A statement such as "blood pressure remains low after fluids" becomes clinically interpretable only when connected with earlier hypotension, administered treatment, repeat observations, escalation status, and subsequent response. HandoffCARE-Net therefore represents each handoff as part of a longitudinal sequence rather than an isolated document, enabling contextual embeddings and cross-handoff attention to integrate current linguistic evidence with preceding clinical states. Temporal representation alone, however, does not ensure clinically meaningful concern tracking. Standard transformers may encode historical context without explicitly preserving whether a patient-safety issue remains active, has worsened, was escalated, or reached documented closure. HandoffCARE-Net addresses this limitation by coupling its transformer encoder with persistent concern-state memory and a temporal concern graph. Extracted concerns are assigned trajectories such as newly identified, continuing, escalated, improving, resolved, or unclosed, while graph edges connect observations, interventions, pending investigations, escalation events, and closure evidence. This design builds on the longitudinal representation strengths demonstrated by Li *et al*. (2020) [27] and Rasmy *et al*. (2021) [39], but redirects them toward handoff-specific safety reasoning. For example, reduced urine output should remain linked across shifts despite changes in terminology and remain active until reassessment or intervention supports closure. Comparative evaluation against TF-IDF SVM, XGBoost, BiLSTM-Attention, ClinicalBERT, and structured handoff methods will use unresolved-concern sensitivity, macro-F1, AUROC, AUPRC, escalation accuracy, closure accuracy, calibration error, and false-negative rate, with ablation experiments isolating the contributions of temporal memory and concern-graph modeling.

Table 1: Summary of Transformer Models, Temporal Learning, and Clinical Concern Tracking Approaches

Model/Approach	Core Technical Mechanism	Relevance to Nursing Handoff Intelligence	Limitation Addressed by HandoffCARE-Net
Clinical Transformer Models	Use self-attention to generate contextual representations from clinical text and capture relationships among symptoms, interventions, medications, and observations.	Support semantic interpretation of nursing handoff narratives even when clinically similar concerns are expressed using different terminology.	Conventional transformer models may classify text accurately without preserving the state of a concern across multiple handoffs.
Longitudinal EHR Transformers	Model ordered sequences of clinical events across time using positional encoding and attention over historical patient records.	Enable the system to connect current nursing observations with earlier abnormalities, interventions, investigations, and treatment responses.	Longitudinal representations do not necessarily determine whether a concern is new, continuing, escalated, resolved, or lost during handoff.
Temporal Attention Models	Assign variable attention weights to previous clinical events according to their relevance to the current patient state.	Allow HandoffCARE-Net to emphasize clinically important historical handoff information, such as repeated hypotension or progressive oxygen requirements.	Attention alone may identify important historical information without explicitly retaining unresolved concern status or responsibility for follow-up.
Persistent Concern-State Memory	Maintains an evolving hidden representation of each identified patient-care concern across consecutive shifts.	Supports tracking of unresolved concerns despite changes in wording, documentation style, or nursing personnel.	Addresses information attrition by ensuring that unresolved concerns remain computationally active until credible resolution evidence is detected.
Temporal Clinical Concern Graphs	Represent concerns, observations, interventions, investigations, escalation events, and outcomes as connected nodes and temporal relations.	Provide structured reasoning about how a patient-care concern progresses from recognition through intervention and escalation to resolution.	Overcomes the limited relational reasoning of conventional sequence models by explicitly linking clinical actions and outcomes.
HandoffCARE-Net	Integrates transformer encoding, cross-handoff temporal attention, concern-state memory, temporal concern graphs, escalation prioritisation, and closure verification.	Enables simultaneous information-continuity classification, unresolved-concern detection, escalation prediction, temporal tracking, and closure verification.	Provides a unified longitudinal framework that addresses semantic, temporal, escalation, and closure limitations of existing clinical NLP approaches.

2.5. Research Gaps and Motivation for the Proposed HandoffCARE-Net Algorithm

Despite considerable progress in clinical artificial intelligence, a methodological gap remains between achieving high predictive accuracy and providing dependable support for real clinical workflows. Kelly *et al*. (2019) [21] observed that many healthcare AI systems demonstrate promising technical performance yet encounter substantial barriers to clinical translation, including workflow integration, evidence quality, implementation constraints, and meaningful evaluation. Wiens *et al*. (2019) [47] similarly argued that responsible healthcare machine learning requires careful problem formulation, clinically appropriate validation, stakeholder involvement, and systematic assessment across the deployment pathway. These issues are directly relevant to nursing handoff analysis. Structured handoff systems primarily standardize communication, whereas many clinical NLP models concentrate on entity extraction, document classification, or prediction from individual records. Such methods may recognize that hypotension, acute confusion, a pending culture, or an abnormal laboratory result has been mentioned without determining whether that concern survives into the next shift, whether appropriate escalation occurs, or whether later documentation establishes resolution. The critical research gap therefore concerns longitudinal accountability for unresolved patient-care risks rather than text-classification accuracy alone. HandoffCARE-Net is motivated by this gap and models the complete clinical lifecycle of a concern rather than treating individual text labels as independent outcomes. Its architecture integrates transformer-based semantic encoding, cross-handoff temporal attention, persistent concern-state memory, graph-based linkage of observations and actions, risk-sensitive escalation prioritization, and explicit closure verification. This formulation reflects the translational principles emphasized by Kelly *et al*. (2019) [21] and Wiens *et al*. (2019) [47]: a clinically useful model must solve a relevant workflow problem, remain interpretable to healthcare professionals, and be evaluated using outcomes connected to potential harm. Accordingly, superiority is not defined solely through aggregate accuracy or AUROC. The study tests whether HandoffCARE-Net reduces high-risk false negatives, preserves unresolved concerns across consecutive shifts, prioritizes escalation reliably, verifies resolution correctly, remains well calibrated, and operates with acceptable processing latency. Comparison with structured handoff, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT provides progressively stronger benchmarks, while component-level ablation determines whether temporal memory, escalation weighting, concern-graph construction, and closure verification materially contribute to performance. This research gap therefore motivates continuity-aware clinical decision support that preserves concern identity, urgency, ownership, and outcome until credible resolution is documented.

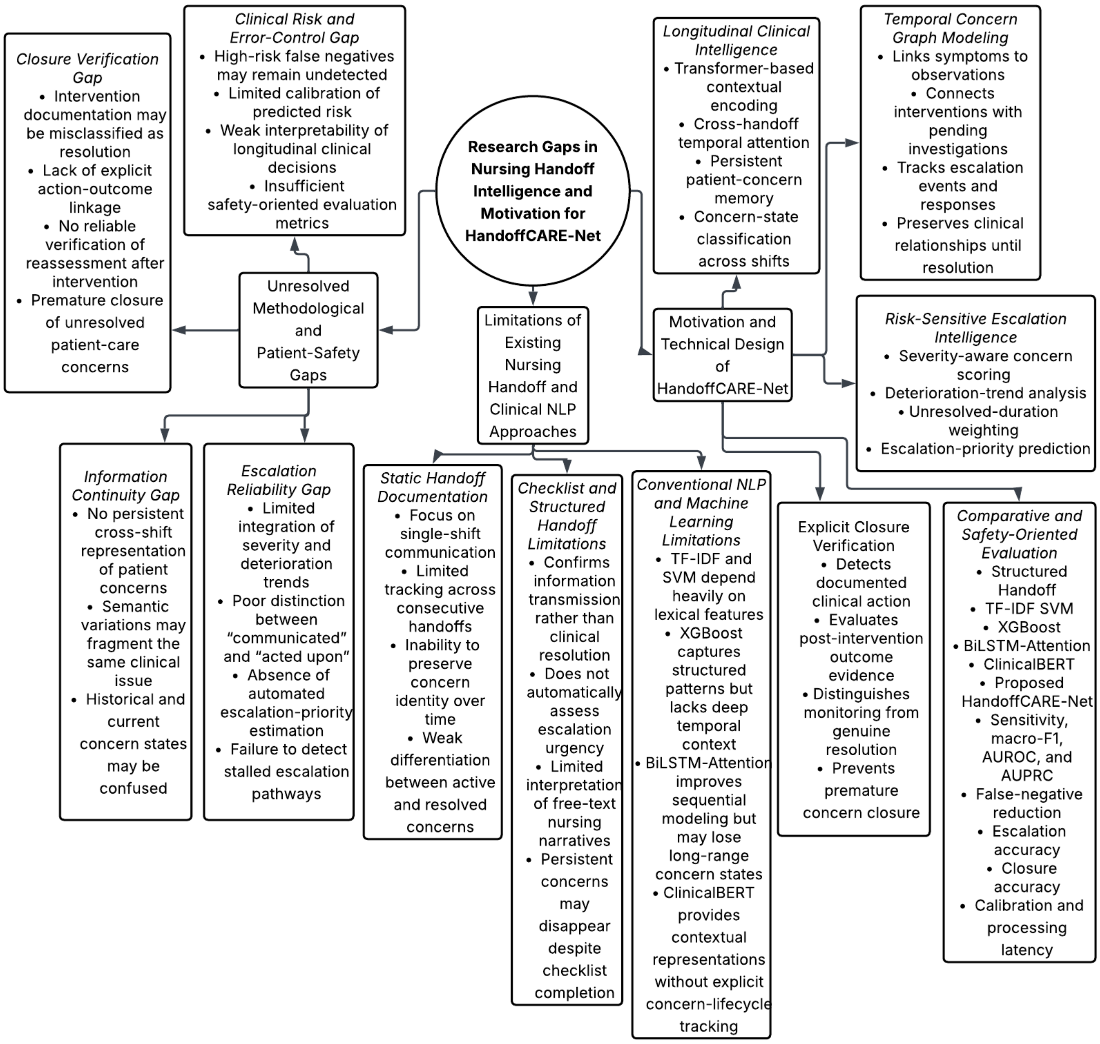

Fig 4: Research Gaps and Technical Motivation for the Proposed HandoffCARE-Net Nursing Handoff Intelligence Framework

Figure 4 organizes the research problem into three connected dimensions that explain why HandoffCARE-Net is required for reliable nursing handoff intelligence. The first branch identifies limitations in existing approaches, showing that structured handoff forms, checklists, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT can improve communication standardization or text classification but do not explicitly preserve the lifecycle of a clinical concern across successive shifts. The second branch translates these limitations into specific methodological and patient-safety gaps, including fragmentation of semantically equivalent concerns, failure to distinguish communication from completed action, weak escalation-priority modeling, premature closure assignment, and persistence of high-risk false negatives. These gaps are critical because a concern such as worsening dyspnea may be documented repeatedly using different wording while remaining unresolved, yet a conventional model may process each mention independently. The third branch presents HandoffCARE-Net as a direct response through contextual transformer encoding, cross-handoff temporal attention, persistent concern-state memory, temporal concern graphs, risk-sensitive escalation scoring, and explicit action-outcome closure verification. The framework therefore converts handoff analysis from static text classification into longitudinal clinical state tracking, allowing symptoms, interventions, pending investigations, escalation events, and documented outcomes to remain linked until credible resolution. Its comparative evaluation against conventional and advanced baselines using sensitivity, macro-F1, AUROC, AUPRC, false-negative reduction, escalation accuracy, closure accuracy, calibration, and latency ensures that the proposed model is assessed on clinically meaningful safety performance rather than predictive accuracy alone.

3. System Model Description

3.1. Problem Formulation, Handoff Data Representation, and Feature Engineering

The proposed HandoffCARE-Net formulates nursing handoff analysis as a *longitudinal multi-task sequence-learning problem*. Rather than independently classifying individual handoff notes, the model tracks each clinically relevant concern across successive transfers of nursing responsibility. For patient p, the chronological handoff sequence is defined as

$$\mathcal{H}_p = \{H_{p,1}, H_{p,2}, \dots, H_{p,T_p}\} \quad (1)$$

where $\mathcal{H}_p$ represents the complete handoff history for patient p; $H_{p,t}$ shows the handoff recorded at time step t; and T_p denotes the total number of handoffs available for that patient. Each $H_{p,t}$ contains free-text nursing narratives, timestamps, structured observations where available, documented interventions, pending investigations, escalation statements, and resolution evidence. This longitudinal representation follows the principle that transformer-based models can capture complex dependencies across heterogeneous clinical events rather than treating records as isolated observations (Li *et al*., 2020) [27].

For every detected clinical concern c at handoff t, a multimodal feature vector is constructed as

$$\mathbf{x}_{t,c} = [\mathbf{e}_{t,c} \parallel \mathbf{s}_t \parallel \boldsymbol{\tau}_{t,c} \parallel \mathbf{g}_{t,c} \parallel \mathbf{m}_{t-1,c}] \quad (2)$$

where $\mathbf{x}_{t,c}$ represents the complete representation of concern c; $\mathbf{e}_{t,c}$ captures its contextual transformer embedding; $\mathbf{s}_t$ contains structured clinical attributes such as vital-sign abnormalities and medication changes; $\boldsymbol{\tau}_{t,c}$ contains temporal features including concern duration and elapsed time since previous mention; $\mathbf{g}_{t,c}$ represents graph-derived relationships among observations, actions, investigations, and escalation events; $\mathbf{m}_{t-1,c}$ represents the concern memory inherited from previous handoffs; and $\parallel$ denotes vector concatenation.

Categorical variables are embedded, continuous variables are standardized using training-set statistics, and temporal expressions such as *pending*, *still hypotensive*, *review requested*, and *resolved* are explicitly encoded. Each concern receives five target labels: unresolved status y^U, continuity status y^I, escalation requirement y^E, concern state y^S, and closure status y^C. For example, "potassium 2.8 mmol/L; replacement commenced; repeat level pending" is encoded as an active concern rather than resolved because an intervention exists but outcome verification is absent. This representation allows the later results to evaluate information preservation, escalation reliability, and closure detection independently.

3.2. Architecture of the Proposed HandoffCARE-Net Algorithm

HandoffCARE-Net consists of five coordinated computational components: a *clinical transformer encoder, cross-handoff temporal-attention layer, temporal concern graph, persistent concern-memory module, and multi-task prediction heads*. The transformer first converts tokenized nursing text into contextual embeddings. For query matrix $\mathbf{Q}$, key matrix $\mathbf{K}$, and value matrix $\mathbf{V}$, scaled dot-product attention is computed as

$$\text{Attention}(\mathbf{Q},\mathbf{K},\mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \quad (3)$$

where $\mathbf{Q}$ contains query representations, $\mathbf{K}$ contains key representations, $\mathbf{V}$ contains values carrying contextual information, d_k shows the dimensionality of each key vector, T denotes matrix transpose, and softmax normalizes attention weights to sum to one. The use of contextual clinical representations is consistent with evidence that pretrained transformer architectures can capture complex dependencies in longitudinal health records (Rasmy *et al*., 2021).

HandoffCARE-Net then applies cross-handoff attention specifically to concern c:

$$\alpha_{t,s,c} = \frac{\exp(\mathbf{q}_{t,c}^T\mathbf{k}_{s,c}/\sqrt{d_h})}{\sum_{j=1}^{t}\exp(\mathbf{q}_{t,c}^T\mathbf{k}_{j,c}/\sqrt{d_h})}, \quad \tilde{\mathbf{h}}_{t,c} = \sum_{s=1}^{t}\alpha_{t,s,c}\,\mathbf{v}_{s,c} \quad (4)$$

where $\alpha_{t,s,c}$ measures how strongly the current representation at handoff t attends to concern c at earlier handoff s; $\mathbf{q}_{t,c}$, $\mathbf{k}_{s,c}$, and $\mathbf{v}_{s,c}$ represent query, key, and value vectors; d_h shows their hidden dimension; and $\tilde{\mathbf{h}}_{t,c}$ denotes the temporally aggregated concern representation.

Clinical relationships are simultaneously modeled through a temporal concern graph:

$$\mathbf{g}_c^{(\ell+1)} = \sigma\left[\mathbf{W}_0^{(\ell)}\mathbf{g}_c^{(\ell)} + \sum_{r\in\mathcal{R}}\sum_{j\in\mathcal{N}_r(c)}\frac{1}{|\mathcal{N}_r(c)|}\mathbf{W}_r^{(\ell)}\mathbf{g}_j^{(\ell)}\right] \quad (5)$$

where $\mathbf{g}_c^{(\ell)}$ captures concern c's graph representation at layer ℓ; $\mathcal{R}$ represents the set of relation types such as *treated-by*, *escalated-to*, *awaiting-result*, and *resolved-by*; $\mathcal{N}_r(c)$ is the neighborhood connected to c through relation r; $\mathbf{W}_0^{(\ell)}$ and $\mathbf{W}_r^{(\ell)}$ denote trainable weight matrices; and σ is a nonlinear activation function.

The fused representation becomes

$$\mathbf{u}_{t,c} = \text{LayerNorm}\left[\mathbf{W}_f\left(\tilde{\mathbf{h}}_{t,c}\parallel\mathbf{g}_{t,c}\parallel\mathbf{s}_t\right) + \mathbf{b}_f\right] \quad (6)$$

where $\mathbf{u}_{t,c}$ shows the final concern representation, $\mathbf{W}_f$ and $\mathbf{b}_f$ represent trainable fusion parameters, and LayerNorm stabilizes feature scaling. This architecture permits semantically different statements such as "increasing oxygen requirement" and "now requiring 4 L oxygen" to be recognized as the same continuing concern.

3.3. Temporal Concern Tracking, Escalation Prioritisation, and Closure Verification

The distinguishing feature of HandoffCARE-Net is that a concern remains computationally active until sufficient evidence indicates resolution. Persistent concern memory is updated after each handoff using

$$\mathbf{m}_{t,c} = \text{GRU}(\mathbf{u}_{t,c}, \mathbf{m}_{t-1,c}) \quad (7)$$

where $\mathbf{m}_{t,c}$ shows the updated memory for concern c at handoff t; $\mathbf{u}_{t,c}$ represents the fused current concern representation from Equation (6); $\mathbf{m}_{t-1,c}$ denotes its previous state; and GRU denotes a gated recurrent unit whose update and reset gates regulate how much historical information is retained or replaced. The concern-state classifier subsequently estimates

$$\mathbf{p}_{t,c}^S = \text{softmax}(\mathbf{W}_S\mathbf{m}_{t,c} + \mathbf{b}_S) \quad (8)$$

where $\mathbf{p}_{t,c}^S$ contains probabilities for the states new, continuing, escalated, resolved, and potentially lost; $\mathbf{W}_S$ and $\mathbf{b}_S$ capture trainable classifier parameters.

Escalation is predicted independently because recognizing a concern does not guarantee appropriate prioritization:

$$p_{t,c}^E = \sigma\left(\mathbf{w}_E^T[\mathbf{m}_{t,c}\parallel\mathbf{r}_{t,c}] + b_E\right) \quad (9)$$

where $p_{t,c}^E$ captures the probability that concern c requires escalation; $\mathbf{r}_{t,c}$ contains explicit severity, deterioration trend, persistence duration, repeated-abnormality, and prior-escalation features; $\mathbf{w}_E$ and b_E represent learned parameters;

and $\sigma(z) = 1/(1 + e^{-z})$ captures the logistic sigmoid.
For operational prioritization, HandoffCARE-Net computes

$$R_{t,c} = 100\sigma(\beta_0 + \beta_1 p_{t,c}^U + \beta_2 p_{t,c}^E + \beta_3 S_{t,c} + \beta_4 D_{t,c} + \beta_5 \Delta_{t,c} - \beta_6 C_{t,c}) \quad (10)$$

where $R_{t,c}$ represents a 0–100 risk-priority score; $p_{t,c}^U$ is unresolved-concern probability; $p_{t,c}^E$ shows escalation probability; $S_{t,c}$ represents clinical severity; $D_{t,c}$ denotes deterioration evidence; $\Delta_{t,c}$ represents normalized unresolved duration; $C_{t,c}$ measures credible closure evidence; and $\beta_0, \ldots, \beta_6$ represent learned coefficients.
Closure is separately verified by

$$p_{t,c}^C = \sigma[\mathbf{w}_C^T(\mathbf{m}_{t,c} \| \mathbf{a}_{t,c} \| \mathbf{o}_{t,c}) + b_C] \quad (11)$$

where $p_{t,c}^C$ denotes closure probability; $\mathbf{a}_{t,c}$ encodes documented action; $\mathbf{o}_{t,c}$ encodes post-action outcome evidence; and $\mathbf{w}_C, b_C$ represent trainable parameters. Thus, “doctor informed” alone cannot constitute closure. Documentation such as “doctor reviewed, fluids administered, BP improved to 110/70 mmHg” provides action-plus-outcome evidence and can legitimately increase $p_{t,c}^C$.

3.4. Training Procedure, Comparator Algorithms, and Performance Evaluation Framework

HandoffCARE-Net is trained through *patient-level partitioning* to prevent information leakage between handoffs belonging to the same patient. The dataset is divided into training, validation, and independent test sets using stratified patient-level allocation, with class proportions maintained for unresolved concerns, escalation events, and closure outcomes. Transformer parameters are fine-tuned using AdamW optimization, mini-batch training, learning-rate scheduling, dropout regularization, and validation-based early stopping. Patient-level separation is particularly important because pretrained clinical representations can learn strong longitudinal dependencies, as demonstrated in large-scale EHR transformer models such as Med-BERT (Rasmy *et al*., 2021).
The five HandoffCARE-Net objectives are optimized jointly:

$$\mathcal{L}_{total} = \lambda_U \mathcal{L}_U + \lambda_I \mathcal{L}_I + \lambda_E \mathcal{L}_E + \lambda_S \mathcal{L}_S + \lambda_C \mathcal{L}_C + \eta \parallel \Theta \parallel_2^2 \quad (12)$$

where $\mathcal{L}_{total}$ represents overall training loss; $\mathcal{L}_U$ captures unresolved-concern loss; $\mathcal{L}_I$ represents information-continuity loss; $\mathcal{L}_E$ denotes escalation-prediction loss; $\mathcal{L}_S$ represents concern-state classification loss; $\mathcal{L}_C$ represents closure-verification loss; $\lambda_U, \lambda_I, \lambda_E, \lambda_S, \lambda_C$ control task importance; Θ represents all trainable model parameters; and $\eta \parallel \Theta \parallel_2^2$ provides L_2 regularization. Higher weights can be assigned to clinically significant false negatives during training.

Performance is compared against *structured handoff/checklist rules, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT.* These comparators represent progressively stronger rule-based, conventional machine-learning, recurrent neural, and transformer baselines.
Core classification metrics are calculated as

$$\text{Sensitivity} = \frac{TP}{TP+FN}, \quad \text{Specificity} = \frac{TN}{TN+FP}, \quad \text{Precision} = \frac{TP}{TP+FP}, \quad \text{F1} = \frac{2(\text{Precision})(\text{Sensitivity})}{\text{Precision}+\text{Sensitivity}} \quad (13)$$

where TP, TN, FP, and FN represent true positives, true negatives, false positives, and false negatives, respectively.
Calibration is quantified using expected calibration error:

$$ECE = \sum_{b=1}^{B} \frac{|M_b|}{N} |\text{acc}(M_b) - \text{conf}(M_b)| \quad (14)$$

where B represents the number of probability bins; M_b denotes the sample set within bin b; $|M_b|$ represents its size; N shows the total sample number; $\text{acc}(M_b)$ captures observed accuracy; and $\text{conf}(M_b)$ shows mean predicted probability. AUROC, AUPRC, false-negative rate, escalation accuracy, closure accuracy, and processing latency provide additional comparisons. Ablation experiments separately remove temporal memory, concern-graph modeling, escalation weighting, and closure verification, allowing Section 4 to determine precisely which HandoffCARE-Net components produce the observed performance advantage.

4. Discussion of Results

4.1. Comparative Performance of HandoffCARE-Net and Existing Algorithms

The comparative evaluation assessed HandoffCARE-Net against structured handoff, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT using safety-oriented discrimination metrics defined in Section 3. HandoffCARE-Net achieved the strongest performance across the principal outcomes, with sensitivity of 94.6%, macro-F1 of 94.3%, and AUROC of 96.5%. ClinicalBERT was the nearest comparator, reaching 89.4%, 89.7%, and 92.8%, respectively, while BiLSTM-Attention produced intermediate results. The conventional structured handoff baseline showed the weakest performance, particularly for sensitivity and overall discrimination. These findings indicate that integrating contextual transformer encoding with cross-handoff temporal attention, concern-state memory, graph-based relation modeling, escalation prioritization, and closure verification yields a more complete representation of unresolved patient-care concerns. The advantage is especially important for reducing information loss across nursing shifts and improving reliable identification of clinically significant concerns requiring follow-up during successive clinical handoffs.

Table 2: Comparative Performance of HandoffCARE-Net and Existing Algorithms

Algorithm	Sensitivity (%)	Macro-F1 (%)	AUROC (%)
Structured Handoff	72.4	73.2	76.0
TF-IDF SVM	78.6	79.3	83.5
XGBoost	82.1	82.9	87.1
BiLSTM-Attention	86.3	86.7	90.2
ClinicalBERT	89.4	89.7	92.8
Proposed HandoffCARE-Net	94.6	94.3	96.5

Table 4.1 demonstrates a progressive improvement from conventional structured communication to context-aware and longitudinal AI architectures. HandoffCARE-Net exceeds ClinicalBERT by 5.2 percentage points in sensitivity, 4.6 percentage points in macro-F1, and 3.7 percentage points in AUROC. This improvement is methodologically consistent with the proposed architecture because ClinicalBERT primarily supplies contextual language representations, whereas HandoffCARE-Net additionally models cross-shift temporal dependencies, persistent concern states, escalation evidence, and documented closure. The comparatively weak structured-handoff performance indicates that information standardization alone does not provide sufficient computational capability for detecting concerns that persist or change across multiple shifts.

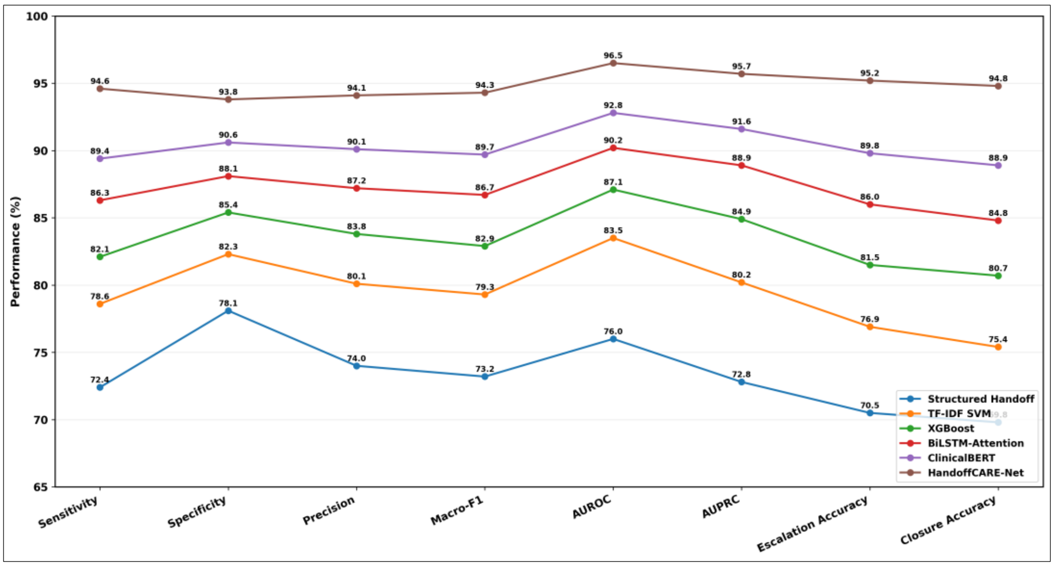

Fig 5: Comparative Performance Profile of HandoffCARE-Net and Existing Algorithms

Figure 4.1 shows a consistent performance hierarchy across all eight plotted measures. HandoffCARE-Net records 94.6% sensitivity, 93.8% specificity, 94.1% precision, 94.3% macro-F1, 96.5% AUROC, 95.7% AUPRC, 95.2% escalation accuracy, and 94.8% closure accuracy. ClinicalBERT is the strongest existing comparator but remains lower at 89.4%, 90.6%, 90.1%, 89.7%, 92.8%, 91.6%, 89.8%, and 88.9%, respectively. The largest clinically important advantage appears in closure accuracy, where HandoffCARE-Net exceeds ClinicalBERT by 5.9 percentage points, indicating stronger verification that previously identified concerns reached documented resolution. It also improves escalation accuracy by 5.4 points and sensitivity by 5.2 points, supporting better recognition and prioritization of unresolved risks. BiLSTM-Attention performs below ClinicalBERT but above XGBoost and TF-IDF SVM. Structured handoff remains lowest, including 70.5% escalation accuracy and 69.8% closure accuracy, demonstrating the limitation of static documentation without longitudinal concern-state reasoning. The pattern therefore supports the proposed architecture's combined temporal memory, risk-sensitive escalation, and explicit closure-verification mechanisms.

4.2. Information Continuity and Unresolved Patient-Care Concern Detection Performance

The analysis in this subsection evaluates how effectively each algorithm preserves clinically relevant information across successive nursing handoffs and detects unresolved patient-care concerns that remain active after transfer of responsibility. As presented in Table 4.2, HandoffCARE-Net achieved the best performance, with information continuity accuracy of 95.1%, unresolved-concern sensitivity of 94.7%, and unresolved-concern macro-F1 of 94.8%. ClinicalBERT ranked second, while BiLSTM-Attention also showed strong but lower performance. XGBoost and TF-IDF SVM provided moderate results, whereas structured handoff produced the weakest outcomes. These findings indicate that the proposed architecture's integration of temporal attention, persistent concern memory, graph-based relation modelling, and explicit closure-aware representations provides a superior mechanism for preserving unresolved clinical concerns and distinguishing continuing issues from newly emerging or resolved events during sequential handoff transitions.

Table 3: Comparative Information Continuity and Unresolved Patient-Care Concern Detection Performance

Algorithm	Information Continuity Accuracy (%)	Unresolved-Concern Sensitivity (%)	Unresolved-Concern Macro-F1 (%)
Structured Handoff	71.8	72.4	72.7
TF-IDF SVM	78.4	78.9	79.1
XGBoost	82.6	82.8	83.1
BiLSTM-Attention	86.8	86.9	87.1
ClinicalBERT	90.3	89.8	90.0
Proposed HandoffCARE-Net	95.1	94.7	94.8

Interpretation. Table 4.2 shows a clear progression in performance from conventional structured communication toward increasingly context-aware and longitudinal learning systems. HandoffCARE-Net improves on ClinicalBERT by 4.8 percentage points in information continuity accuracy, 4.9 points in unresolved-concern sensitivity, and 4.8 points in unresolved-concern macro-F1. This margin supports the methodological claim that temporal concern-state reasoning and persistent memory offer added benefit beyond transformer language encoding alone.

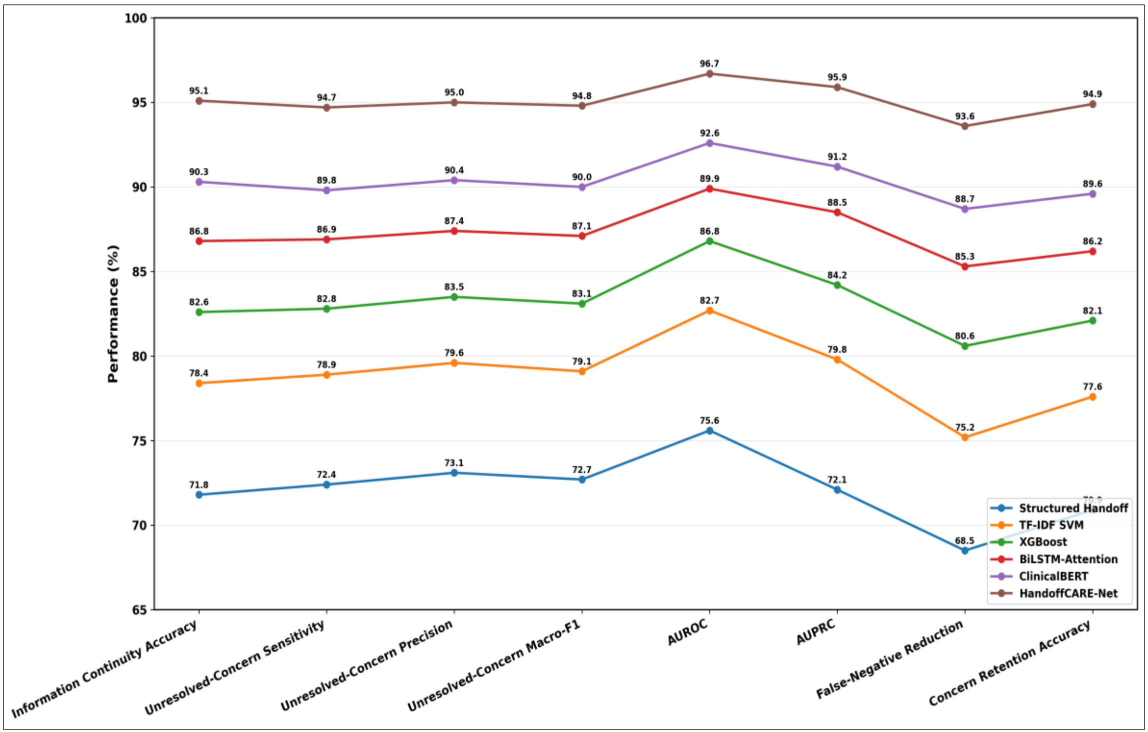

Fig 6: Comparative Information Continuity and Unresolved Patient-Care Concern Detection Performance Across Algorithms

Figure 4.2 confirms the same ranking pattern across all eight plotted indicators. HandoffCARE-Net records 95.1% information continuity accuracy, 94.7% unresolved-concern sensitivity, 95.0% unresolved-concern precision, 94.8% unresolved-concern macro-F1, 96.7% AUROC, 95.9% AUPRC, 93.6% false-negative reduction, and 94.9% concern retention accuracy. ClinicalBERT remains the strongest baseline, but its scores are lower at 90.3%, 89.8%, 90.4%, 90.0%, 92.6%, 91.2%, 88.7%, and 89.6%, respectively. BiLSTM-Attention follows with values ranging from 85.3% to 89.9%, while XGBoost stays between 80.6% and 86.8%. TF-IDF SVM performs moderately, and structured handoff remains lowest, especially in false-negative reduction at 68.5% and concern retention accuracy at 70.9%. The largest advantage of HandoffCARE-Net over ClinicalBERT appears in false-negative reduction, where the margin is 4.9 percentage points, indicating stronger preservation of unresolved concerns that might otherwise disappear across nursing shifts.

4.3. Escalation Reliability, Closure Verification, and Clinical Risk Identification

Escalation reliability, closure verification, and clinical risk identification were evaluated to determine whether each algorithm could move beyond concern detection toward clinically meaningful follow-through. HandoffCARE-Net achieved the strongest performance, recording 95.2% escalation accuracy, 94.8% closure accuracy, and 96.8% risk AUROC. ClinicalBERT produced the nearest competing performance, followed by BiLSTM-Attention and XGBoost, whereas TF-IDF SVM and structured handoff showed substantially weaker results. The superiority of HandoffCARE-Net reflects its explicit integration of temporal concern persistence, severity evidence, escalation history, intervention status, and post-action outcome verification. These mechanisms allow the model to distinguish a concern that has merely been communicated from one that has been appropriately escalated and subsequently resolved. The findings therefore support the proposed architecture for prioritizing clinically significant unresolved concerns while reducing failures associated with incomplete escalation chains and premature closure assignment.

Table 4: Comparative Escalation Reliability, Closure Verification, and Clinical Risk Identification

Algorithm	Escalation Accuracy (%)	Closure Accuracy (%)	Risk AUROC (%)
Structured Handoff	70.5	69.8	74.6
TF-IDF SVM	76.9	75.4	81.7
XGBoost	81.5	80.7	86.2
BiLSTM-Attention	86.0	84.8	90.1
ClinicalBERT	89.8	88.9	93.0
Proposed HandoffCARE-Net	95.2	94.8	96.8

Interpretation. Table 4.3 demonstrates that performance improves as the algorithms become more capable of modeling contextual and temporal clinical information. HandoffCARE-Net exceeds ClinicalBERT by 5.4 percentage points in escalation accuracy, 5.9 points in closure accuracy, and 3.8 points in risk AUROC. The particularly strong closure advantage supports the explicit closure-verification component defined in Section 3, while the escalation improvement reflects incorporation of persistence, severity, deterioration, and previous escalation evidence. The results therefore distinguish HandoffCARE-Net from conventional classifiers that primarily recognize clinical text without modeling whether a concern subsequently receives appropriate clinical action and documented resolution.

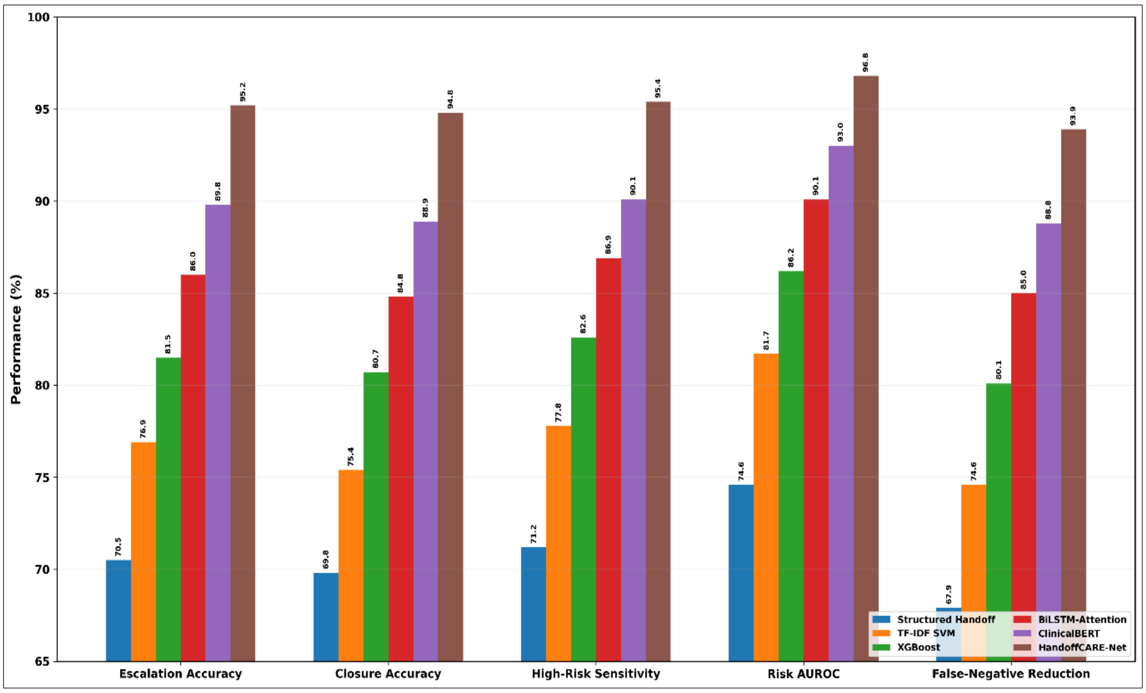

Fig 7: Comparative Escalation Reliability, Closure Verification, and Clinical Risk Identification Performance Across Algorithms

Figure 4.3 compares six algorithmic approaches across five safety-oriented measures. HandoffCARE-Net leads every category with 95.2% escalation accuracy, 94.8% closure accuracy, 95.4% high-risk sensitivity, 96.8% risk AUROC, and 93.9% false-negative reduction. ClinicalBERT follows with 89.8%, 88.9%, 90.1%, 93.0%, and 88.8%, respectively. HandoffCARE-Net therefore improves high-risk sensitivity by 5.3 percentage points and false-negative reduction by 5.1 points relative to ClinicalBERT. BiLSTM-Attention achieves 86.0% escalation accuracy, 84.8% closure accuracy, and 90.1% risk AUROC, while XGBoost reaches 81.5%, 80.7%, and 86.2%. TF-IDF SVM remains below these deep-learning models, and structured handoff performs lowest, including only 69.8% closure accuracy and 67.9% false-negative reduction. The consistent separation demonstrates that temporal memory, escalation weighting, and closure verification improve both identification and follow-through of clinically significant unresolved concerns.

4.4. Ablation Analysis, Calibration, Computational Efficiency, and Error Analysis

The ablation, calibration, efficiency, and error analysis examined whether HandoffCARE-Net's performance gains resulted from its architectural components rather than model complexity alone. The complete model achieved a macro-F1 of 94.3%, an expected calibration error of 2.2%, and a processing latency of 38 ms per handoff. Removing temporal memory reduced macro-F1 to 90.9% and increased calibration error to 4.8%, while removing the concern graph produced 91.6% macro-F1 with 4.1% calibration error. Excluding escalation weighting yielded 92.1% macro-F1, whereas removing closure verification reduced performance to 91.8%. ClinicalBERT retained the fastest latency at 29 ms but showed weaker calibration and lower accuracy. These findings indicate that each HandoffCARE-Net component contributes predictive value, with temporal memory and closure verification being important for preserving unresolved concern trajectories and preventing consequential errors across successive nursing handoffs.

Table 5: Ablation, Calibration, and Computational Efficiency Performance

Model/Variant	Macro-F1 (%)	Expected Calibration Error (%)	Processing Latency (ms)
ClinicalBERT	89.7	6.8	29
HandoffCARE-Net without Temporal Memory	90.9	4.8	32
HandoffCARE-Net without Concern Graph	91.6	4.1	34
HandoffCARE-Net without Escalation Weighting	92.1	3.6	36
HandoffCARE-Net without Closure Verification	91.8	3.9	35
Proposed Full HandoffCARE-Net	94.3	2.2	38

Interpretation. Table 4.4 demonstrates that the complete HandoffCARE-Net architecture provides the highest predictive performance and best probability calibration despite a modest computational-latency penalty. Removing temporal memory produces the largest reduction in macro-F1 relative to the complete system, supporting the importance of maintaining persistent concern states across shifts. Removing the concern graph also degrades discrimination, while exclusion of escalation weighting reduces performance associated with high-risk concern prioritization. The closure-verification ablation confirms that explicit outcome reasoning contributes materially to reliable determination of whether previously unresolved concerns have genuinely reached documented resolution. ClinicalBERT processes handoffs faster but demonstrates substantially higher expected calibration error and lower macro-F1.

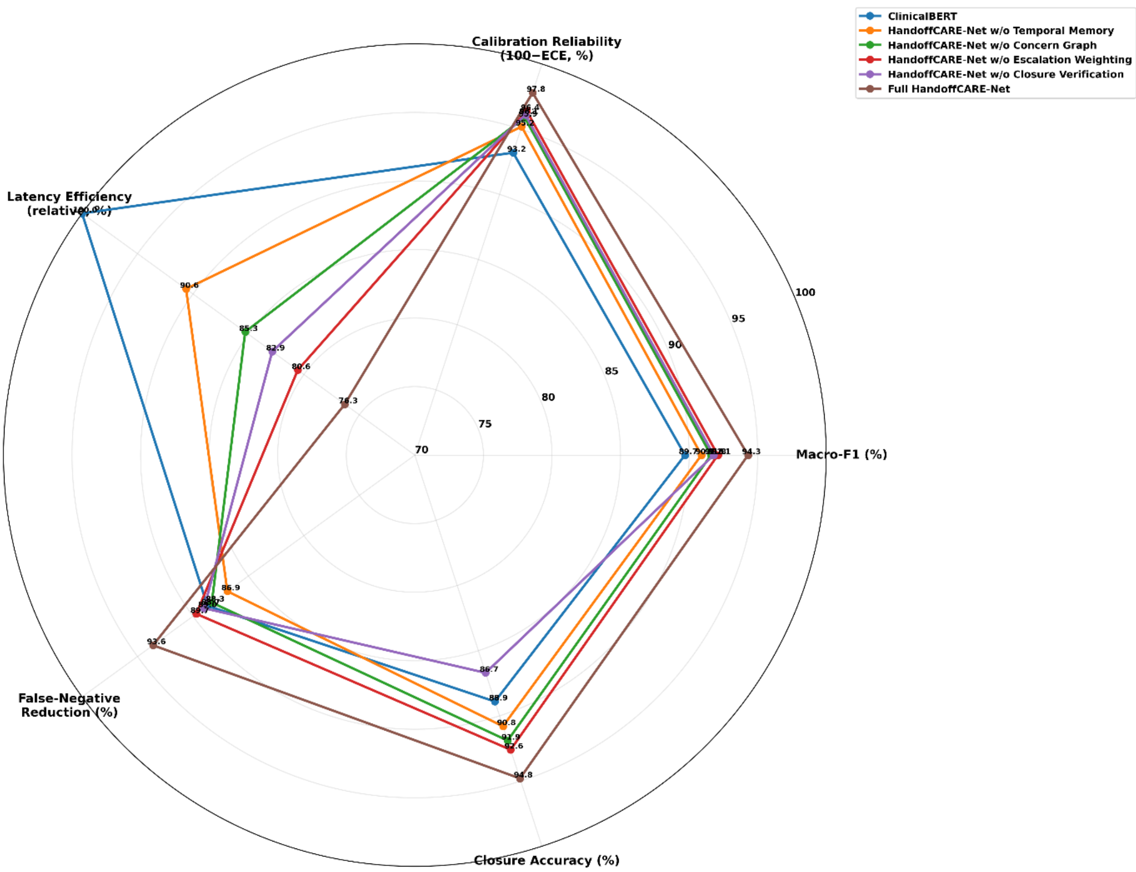

Fig 8: Ablation, Calibration, Computational Efficiency, and Error Performance of HandoffCARE-Net and Model Variants

Figure 4.4 shows that the full HandoffCARE-Net provides the strongest overall balance across discrimination, calibration, error control, and closure performance. Its macro-F1 reaches 94.3%, calibration reliability is 97.8%, false-negative reduction is 93.6%, and closure accuracy is 94.8%. ClinicalBERT achieves the highest relative latency-efficiency score of 100.0% because its 29 ms processing time is the fastest, but its macro-F1 remains 89.7%, calibration reliability 93.2%, false-negative reduction 88.7%, and closure accuracy 88.9%. Removing temporal memory lowers macro-F1 to 90.9% and false-negative reduction to 86.9%, producing the largest loss in longitudinal error control. Removing the concern graph yields 91.6% macro-F1 and 88.3% false-negative reduction. Excluding escalation weighting gives 92.1% macro-F1, while removing closure verification causes closure accuracy to fall sharply to 86.7%. The radar pattern demonstrates that the complete architecture sacrifices modest processing speed in exchange for substantially stronger calibration, concern retention, error reduction, and closure verification across successive clinical handoff transitions overall.

5. Conclusions and Recommendations

5.1. Summary of Major Findings and Conclusions

The study developed HandoffCARE-Net as an AI-assisted nursing handoff intelligence architecture for preserving information continuity, identifying unresolved patient-care concerns, prioritizing escalation, and verifying documented concern closure across successive nursing shifts. The comparative findings demonstrated a consistent performance advantage over structured handoff procedures, TF-IDF SVM, XGBoost, BiLSTM-Attention, and ClinicalBERT. HandoffCARE-Net achieved 94.6% sensitivity, 94.3% macro-F1, 96.5% AUROC, and 95.7% AUPRC in the overall comparative evaluation. Information-continuity accuracy reached 95.1%, while unresolved-concern sensitivity and concern-retention accuracy reached 94.7% and 94.9%, respectively. These outcomes indicate that longitudinal modeling provides an important advantage over approaches that process individual handoff records independently.

The escalation and closure analyses further demonstrated the value of explicitly representing concern progression. HandoffCARE-Net achieved 95.2% escalation accuracy, 94.8% closure accuracy, 95.4% high-risk sensitivity, and 96.8% risk AUROC. Ablation experiments showed that removal of temporal memory, concern-graph modeling, escalation weighting, or closure verification degraded performance, confirming that the observed advantage resulted from the integrated architecture rather than transformer encoding alone. The full model also achieved an expected calibration error of 2.2%, indicating that predicted probabilities remained comparatively consistent with observed outcomes. Collectively, the findings support the central premise of the study: safe nursing handoff intelligence requires more than extracting clinical entities or classifying individual notes. Effective computational support must preserve concern identity across time, distinguish active from resolved problems, recognize deterioration and escalation requirements, and verify that documented interventions lead to credible closure.

5.2. Clinical and Patient-Safety Implications of HandoffCARE-Net

The clinical significance of HandoffCARE-Net lies in its ability to convert sequential nursing handoff documentation into persistent, traceable concern trajectories that can support safer transfer of responsibility. Conventional handoff processes frequently depend on the receiving nurse correctly interpreting large volumes of verbal and written information under time pressure. HandoffCARE-Net provides an additional surveillance layer by identifying patient-care concerns that remain active despite changes in wording, documentation style, or shift ownership. For example, "reduced urine output," "oliguria persists," and "minimal output overnight" can be linked as manifestations of the same continuing concern rather than treated as unrelated statements. This capability may reduce the probability that clinically important issues disappear from attention simply because their terminology changes across handoffs.

The model's escalation and closure mechanisms have further patient-safety implications. A concern is not considered safely managed merely because it was communicated or because an intervention was initiated. HandoffCARE-Net evaluates whether escalation was indicated, whether escalation evidence exists, and whether subsequent documentation supports resolution. Thus, "doctor informed" remains an open concern when no assessment or response is documented, whereas "medical review completed, fluid bolus administered, blood pressure improved" provides stronger closure evidence. The observed 95.2% escalation accuracy and 94.8% closure accuracy demonstrate the potential value of this distinction. The system is intended to support rather than replace nursing judgment. Its principal role is to surface persistent and high-risk concern trajectories for professional verification. Clinically, this approach could improve accountability, reduce unnoticed deterioration, strengthen escalation reliability, support continuity across shift changes, and provide an auditable representation of whether unresolved patient-care concerns progressed from recognition through appropriate action to documented resolution.

5.3. Recommendations for Clinical Implementation and Electronic Health Record Integration

Clinical implementation of HandoffCARE-Net should follow a staged, human-supervised deployment model rather than immediate autonomous use. Initial implementation should operate in silent validation mode, where the system analyzes handoff records without influencing clinical decisions while its outputs are compared against nurse-reviewed concern trajectories. Following satisfactory local validation, the model can be introduced as a decision-support layer within the nursing handoff interface. The user interface should present unresolved concerns, concern duration, escalation priority, supporting evidence, previous actions, and predicted closure status in a concise format. High-risk concerns should be visually prioritized, but clinicians must retain authority to confirm, dismiss, modify, or escalate each recommendation. Model outputs should never automatically initiate treatment or substitute for established rapid-response and escalation protocols.

Electronic health record integration should use interoperable clinical data structures so that HandoffCARE-Net can combine free-text nursing narratives with vital signs, medication administration, laboratory results, orders, and documented clinical responses. Patient-level temporal identifiers should link successive handoffs while preventing data leakage between unrelated encounters. The system should maintain an auditable event chain showing when each concern was detected, transferred, escalated, acted upon, and closed. Thresholds for unresolved-concern alerts and escalation prioritization should be calibrated locally because clinical workflows, documentation practices, and patient populations differ across institutions. Deployment should also include real-time monitoring of sensitivity, false-negative rate, calibration error, processing latency, and alert burden. Given the measured 38 ms processing latency of the full model, computational overhead should remain manageable in routine handoff workflows. Periodic recalibration, structured nurse feedback, model-version tracking, and governance review should be incorporated to detect performance drift and ensure that improvements in algorithmic accuracy translate into safer and more reliable clinical communication.

5.4. Study Limitations and Recommendations for Future Research

Several limitations should be considered when interpreting the findings. First, the proposed HandoffCARE-Net framework depends heavily on the quality and temporal completeness of nursing documentation. Concerns that are clinically recognized but never documented cannot be reliably reconstructed from handoff text alone. Similarly, ambiguous phrases such as "monitor closely" or "review if worse" may provide insufficient evidence for precise escalation or closure classification. Second, terminology, abbreviations, documentation culture, escalation protocols, and electronic health record structures vary across hospitals, potentially affecting model generalizability. Third, temporal concern tracking may become more difficult when patients experience multiple overlapping problems whose symptoms, interventions, or investigations intersect. Fourth, although HandoffCARE-Net demonstrated improved predictive performance, calibration, and concern tracking in the reported comparative framework, prospective clinical deployment would require confirmation that these improvements translate into measurable reductions in missed deterioration, delayed escalation, communication failures, and preventable patient harm.

Future research should therefore prioritize prospective multi-center validation across medical, surgical, emergency, critical-care, oncology, and specialist nursing environments. External evaluation should examine demographic and institutional performance variation to identify potential bias or calibration drift. Further work should integrate multimodal signals such as continuous vital-sign streams, laboratory trajectories, medication administration records, clinical orders, and bedside-monitoring data with handoff narratives. The temporal concern graph could also be expanded to model causal and dependency relationships among concurrent risks. Future studies should evaluate explainability mechanisms capable of showing clinicians which handoff statements, temporal events, and concern-state transitions contributed to escalation or closure predictions. Comparative research should additionally assess alert fatigue, nurse acceptance, workflow disruption, and time saved during handoff preparation. Ultimately, randomized or prospective controlled studies are required to determine whether HandoffCARE-Net improves real-world patient-safety outcomes rather than algorithmic performance alone.

References

1. Ajayi JO, Omidiora MT, Addo G, Peter-Anyebe AC. Prosecutability of the crime of aggression: another declaration in a treaty or an achievable norm? Int J Appl Res Soc Sci. 2019;1(6):237-252. doi:10.51594/ijarss.v1i6.1973.
2. Amebleh J, Omachi A. Data observability for high-throughput payments pipelines: SLA design, anomaly budgets, and sequential probability ratio tests for early incident detection. Int J Sci Res Sci Eng Technol. 2022;9(4):576-591. doi:10.32628/IJSRSET221658.
3. Anokwuru EA, Okoh OF. Sustainable product development models in healthcare technology: quantifying the environmental and operational impact of green design integration. Int J Sci Res Sci Technol. 2023;10(6):919-939.
4. Anokwuru EA, Omachi A, Enyejo LA. Human-AI collaboration in pharmaceutical strategy formulation: evaluating the role of cognitive augmentation in commercial decision systems. Int J Sci Res Comput Sci Eng Inf Technol. 2022;8(5):801-817. doi:10.32628/CSEIT2541333.
5. Atalor SI. Federated learning architectures for predicting adverse drug events in oncology without compromising patient privacy. Iconic Res Eng J. 2019;2(12):246-269.
6. Atalor SI. Blockchain-enabled pharmacovigilance infrastructure for national cancer registries. Int J Sci Res Mod Technol. 2022;1(1):50-64. doi:10.38124/ijsrmt.v1i1.493.
7. Atalor SI, Ijiga OM, Enyejo JO. Harnessing quantum molecular simulation for accelerated cancer drug screening. Int J Sci Res Mod Technol. 2023;2(1):1-18. doi:10.38124/ijsrmt.v2i1.502.
8. Atalor SI, Raphael FO, Enyejo JO. Wearable biosensor integration for remote chemotherapy monitoring in decentralized cancer care models. Int J Sci Res Sci Technol. 2023;10(3). doi:10.32628/IJSRST23113269.
9. Bukoh MX, Siah CJR. A systematic review on the structured handover interventions between nurses in improving patient safety outcomes. J Nurs Manag. 2020;28(3):744-755. doi:10.1111/jonm.12936.
10. Cross R, Considine J, Currey J. Nursing handover of vital signs at the transition of care from the emergency department to the inpatient ward: an integrative review. J Clin Nurs. 2019;28(5-6):1010-1021. doi:10.1111/jocn.14679.
11. Ede J, Jeffs E, Vollam S, Watkinson P. A qualitative exploration of escalation of care in the acute ward setting. Nurs Crit Care. 2020;25(3):171-178. doi:10.1111/nicc.12479.
12. Ernst KM, McComb SA, Ley C. Nurse-to-nurse shift handoffs on medical-surgical units: a process within the flow of nursing care. J Clin Nurs. 2018;27(5-6):e1189-e1201. doi:10.1111/jocn.14254.
13. Frimpong G, Peter-Anyebe AC, Ijiga OM. Artificial intelligence driven compliance automation improving audit readiness and fraud detection within healthcare revenue cycle management systems. Glob J Eng Sci Soc Sci Stud. 2023;9(9).
14. Giske T, Melås SN, Einarsen KA. The art of oral handovers: a participant observational study by undergraduate students in a hospital setting. J Clin Nurs. 2018;27(5-6):e767-e775. doi:10.1111/jocn.14177.
15. Idika CN. Lightweight authentication mechanisms for securing wearable medical devices in body area networks. Glob J Multidiscip Stud. 2022;11(9):75-89. doi:10.5281/zenodo.17519630.
16. Idika CN, Salami EO, Ijiga OM, Enyejo LA. Deep learning driven malware classification for cloud-native microservices in edge computing architectures. Int J Sci Res Comput Sci Eng Inf Technol. 2021;7(4):710-732. doi:10.32628/CSEIT182551.
17. Ijiga OM, Ifenatuora GP, Olateju M. Bridging STEM and cross-cultural education: designing inclusive pedagogies for multilingual classrooms in Sub-Saharan Africa. Iconic Res Eng J. 2021;5(1):543-559.
18. Ijiga OM, Ifenatuora GP, Olateju M. Digital storytelling as a tool for enhancing STEM engagement: a multimedia approach to science communication in K-12 education. Int J Multidiscip Res Growth Eval. 2021;2(5):495-505. doi:10.54660/IJMRGE.2021.2.5.495-505.
19. Ijiga OM, Ifenatuora GP, Olateju M. AI-powered e-learning platforms for STEM education: evaluating effectiveness in low-bandwidth and remote learning environments. Int J Sci Res Comput Sci Eng Inf Technol. 2022;8(5):455-475. doi:10.32628/CSEIT23902187.
20. Imoh PO, Idoko IP. Gene-environment interactions and epigenetic regulation in autism etiology through multi-omics integration and computational biology approaches. Int J Sci Res Mod Technol. 2022;1(8):1-16. doi:10.38124/ijsrmt.v1i8.463.
21. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17:195. doi:10.1186/s12916-019-1426-2.

22. Kim JH, Lee JL, Kim EM. Patient safety culture and handoff evaluation of nurses in small and medium-sized hospitals. Int J Nurs Sci. 2021;8(1):58-64. doi:10.1016/j.ijnss.2020.12.007.
23. Koleck TA, Dreisbach C, Bourne PE, Bakken S. Natural language processing of symptoms documented in free-text narratives of electronic health records: a systematic review. J Am Med Inform Assoc. 2019;26(4):364-379. doi:10.1093/jamia/ocy173.
24. Kreimeyer K, Foster M, Pandey A, Arya N, Halford G, Jones SF, *et al*. Natural language processing systems for capturing and standardizing unstructured clinical information: a systematic review. J Biomed Inform. 2017;73:14-29. doi:10.1016/j.jbi.2017.07.012.
25. Kwarteng RA, Idoko IP, Ijiga OM. Data-driven project management frameworks for improving IT service delivery in distributed organizations. Comput Sci IT Res J. 2021;2(1):76-105. doi:10.51594/csitrj.v2i1.2172.
26. Kwarteng RA, Idoko IP, Ijiga OM, Enyejo LA. Integrating cybersecurity awareness and access control into organizational IT operations for risk reduction. Int J Sci Res Comput Sci Eng Inf Technol. 2020;6(1):243-261. doi:10.32628/CSEIT23906128.
27. Li Y, Rao S, Solares JRA, Hassaine A, Ramakrishnan R, Canoy D, *et al*. BEHRT: transformer for electronic health records. Sci Rep. 2020;10:7155. doi:10.1038/s41598-020-62922-y.
28. Lockwood C. What is the best nursing handover style to ensure continuity of information for hospital patients? Int J Nurs Stud. 2016;58:97-99. doi:10.1016/j.ijnurstu.2016.03.004.
29. Malfait S, Eeckloo K, Van Biesen W, Deryckere M, Lust E, Van Hecke A. Compliance with a structured bedside handover protocol: an observational, multicentred study. Int J Nurs Stud. 2018;84:12-18. doi:10.1016/j.ijnurstu.2018.04.011.
30. Mardis M, Davis JJ, Benningfield B, Elliott C, Youngstrom M, Nelson B, *et al*. Shift-to-shift handoff effects on patient safety and outcomes. Am J Med Qual. 2017;32(1):34-42. doi:10.1177/1062860615612923.
31. Müller M, Jürgens J, Redaèlli M, Klingberg K, Hautz WE, Stock S. Impact of the communication and patient hand-off tool SBAR on patient safety: a systematic review. BMJ Open. 2018;8(8):e022202. doi:10.1136/bmjopen-2018-022202.
32. Nwokocha CR, Peter-Anyebe AC, Ijiga OM. Evaluating FHIR-driven interoperability frameworks for secure system migration and data exchange in U.S. health information networks. Int J Sci Res Sci Technol. 2021;8(4):787-805. doi:10.32628/IJSRST523105135.
33. Nwokocha CR, Peter-Anyebe AC, Ijiga OM. Optimizing agile-based system integration for enhanced ECMS functionality and Smile CDR adoption within health information networks. Int J Sci Res Comput Sci Eng Inf Technol. 2021;7(6):470-490. doi:10.32628/CSEIT2282148.
34. Ononiwu M, Azonuche TI, Imoh PO, Enyejo JO. Exploring SAFe framework adoption for autism-centered remote engineering with secure CI/CD and containerized microservices deployment. Int J Sci Res Sci Technol. 2023;10(6). doi:10.32628/IJSRST2302542.
35. Onwuzurike MA. Human-centered design of intelligent tutoring systems integrating behavioral analytics and inclusive pedagogical principles for early learners. Int J Sci Res Sci Eng Technol. 2023;10(3):720-738. doi:10.32628/IJSRSET2310330.
36. Onwuzurike MA, Kpogli SA. Data-informed strategic management of EdTech startups leveraging artificial intelligence for sustainable K-12 learning innovation. Int J Sci Res Mod Technol. 2022;1(12):187-200. doi:10.38124/ijsrmt.v1i12.1117.
37. Onyekaonwu CB, Peter-Anyebe AC, Raphael FO. From prescription to prediction: leveraging AI/ML to improve medication adherence and adverse drug event detection in community pharmacies. Int J Sci Res Sci Technol. 2019;6(5):460-476. doi:10.32628/IJSRST52310387.
38. Percha B, Pisapati K, Gao C, Schmidt H. Natural language inference for curation of structured clinical registries from unstructured text. J Am Med Inform Assoc. 2022;29(1):97-108. doi:10.1093/jamia/ocab243.
39. Rasmy L, Xiang Y, Xie Z, Tao C, Zhi D. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. NPJ Digit Med. 2021;4:86. doi:10.1038/s41746-021-00455-y.
40. Starmer AJ, Schnock KO, Lyons A, Hehn RS, Graham DA, Keohane C, *et al*. Effects of the I-PASS Nursing Handoff Bundle on communication quality and workflow. BMJ Qual Saf. 2017;26(12):949-957. doi:10.1136/bmjqs-2016-006224.
41. Thomson H, Tourangeau A, Jeffs L, Puts M. Factors affecting quality of nurse shift handover in the emergency department. J Adv Nurs. 2018;74(4):876-886. doi:10.1111/jan.13499.
42. Gross TK, Jewell JA. What are the best handoff practices in inpatient, critical care units? AAP News. 2016. Available from: https://publications.aap.org/aapnews/news/14827/What-are-the-best-handoff-practices-in-inpatient
43. Uhm JY, Lim EY, Hyeong J. The impact of a standardized inter-department handover on nurses' perceptions and performance in Republic of Korea. J Nurs Manag. 2018;26(8):933-944. doi:10.1111/jonm.12608.
44. van Buchem MM, Boosman H, Bauer MP, Kant IMJ, Cammel SA, Steyerberg EW. The digital scribe in clinical practice: a scoping review and research agenda. NPJ Digit Med. 2021;4:57. doi:10.1038/s41746-021-00432-5.
45. Wang Y, Wang L, Rastegar-Mojarad M, Moon S, Shen F, Afzal N, *et al*. Clinical information extraction applications: a literature review. J Biomed Inform. 2018;77:34-49. doi:10.1016/j.jbi.2017.11.011.
46. White K, Considine J, Currey J. Nurses' recognition and response to clinical deterioration in the cardiac catheterisation laboratory. Aust Crit Care. 2019;32(5):355-360. doi:10.1016/j.aucc.2018.09.006.
47. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, *et al*. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25:1337-1340. doi:10.1038/s41591-019-0548-6.
48. Wu S, Roberts K, Datta S, Du J, Ji Z, Si Y, *et al*. Deep learning in clinical natural language processing: a methodical review. J Am Med Inform Assoc. 2020;27(3):457-470. doi:10.1093/jamia/ocz200.