ADROID Conference Series: Engineering and Technology
Volume 2, Issue 2, pp.79-87, 2026

AI-Based Detection of Care Gaps in Electronic Health Records for Improved Clinical Decision Making Using Temporal, Semantic, and Guideline-Aware Learning

Vidhika, Manvi Chahar, Shivani Jaswal, Ayush Mehta

Apex Institute of Technology (CSE), Chandigarh University, Punjab, India

22bai70779@cuchd.in, 22bai70461@cuchd.in, 22bai71290@cuchd.in, 22bai71386@cuchd.in

Abstract. Lapses in care (i.e., missed screenings, late follow-ups, missed medications, off-pathway deviations) remain a significant source of avoidable harm despite the widespread implementation of Electronic Health Records (EHRs) and clinical decision support systems based on rules (Clinical Decision Support, 2009). To address this issue, this manuscript proposes a guideline-aware artificial intelligence system that actively identifies gaps in care by combining structured EHR variables with unstructured clinical documents and tracking patient progress. Clinical guidelines are represented as machine-readable paths in a body of knowledge graph, which explicitly allow one to compare a patient's sequence of temporal events with the care actions desired for it. A hybrid structure integrates the BioBERT-based entity extraction functionality with a Transformer to longitudinally model and an attention-based fusion layer to generate a care-gap risk score with understandable explanations. Benchmark critical-care EHR data experiments expected results are high-precision of the system, a less-fatigued alert, and earlier detection of missed interventions as compared to rule-based CDSS and traditional machine-learned baselines. This designation takes clinical decision support to the next level by offering more than reactive alerts; it provides proactive, explainable detection of guideline violations, making it safer and more effective in providing quality care to patients.

Keywords: Care gap detection, Clinical decision support, electronic health records, Explainable AI, Guideline-aware learning, Temporal deep learning.

1. INTRODUCTION

EHRs have revolutionized the healthcare document as the histories of patients, their laboratory findings, medication intakes, imaging studies and notes of clinicians are converted into digitized long-term electronic profiles. Epic and Cerner (which has since become Oracle Health) are systems which allow clinicians to find incredible amounts of patient information at point of care.[1], [2] However, this has led to an increased clinical decision burden; physicians need to make sense of heterogeneous data streams over time, and integrate conflicting notes, and make sure that clinical guidelines are adhered to as clinical rules change in a short consultation. This cognitive overload is causing forgotten follow-ups, screenings and medication reconciliation mistakes as well as tardiness of interventions all which are defined as care gaps.[3]

A care gap is defined as an inconsistency between the reported care path of a patient and the evidence-based care recommendation (e.g., the missed HbA1c reading in diabetes, the delay in hypertension check-up, the absence of cancer screening, or medication non-adherence). Most of the deployed solutions, even though they integrate CDSS into EHRs, are rules-based alerts that are activated when a person crosses a simple threshold (such as laboratory value cut-offs or set time alarms). Such systems are often affected with alert fatigue issues, low level of personalization, and low awareness of patients context.[4], [5]

A number of technical issues exert this issue. To start with, a significant part of EHR data is as a unstructured clinical note, discharge summary and prescription which are not directly machine-readable.[6] Second, the history of the patients is irregular in time: the visits appear with intervals that are not evenly spaced and changing diagnosis and treatment. Third, individual differences in patients imply that the same rules cannot be consistently used across age factors, comorbidities, or risk factors. Lastly, clinical guidelines provided by government bodies and organizations like the World Health Organization, the National Institute of Health and Care Excellence, and the American Diabetes Association are not straightforward, conditional and hard to program including in fixed rules engines.[7]

How can a guideline-aware AI system effectively detect care gaps in EHR data while improving precision, reducing alert fatigue, and ensuring explainability compared to traditional CDSS?

It is these restrictions that drive the demand towards an artificial intelligence powered, intelligent and proactive structure that is able to be learnt on patient health through EHR, understand textual clinical input and contextually examine the real patient care provided against programmed directions within machine-coded therapy guidelines. Instead of providing reactive warnings, this kind of system would also be capable of detecting the new lapses in care, measure the potential risk, and provide interpretable evidence to clinicians to help them act as fast as possible.[8], [9]

Contributions of the Paper

- Perception-based AI system that projects patient EHR events to formalized clinical criteria out of ADA, WHO, and NICE based on a machine-readable knowledge graph.
- Temporal modelling of patient history: Time-varying clinical evolution can be modeled by sequence learning (Transformer/LSTM) to utilize patient history.
- Structured and unstructured EHR data merging via NLP based entity extraction of clinical notes and coded medical records.
- Shift from prediction process deviation detection.

Explainable clinical reasoning (not black-box).

2. LITERATURE REVIEW

A. Traditional Rule-Based Clinical Decision Support Systems (CDSS).

Other initial uses of CDSS, implemented into an EHR system like Epic System or Cerner, had deterministic rules (e.g., threshold lab values, fixed reminders, drug–drug interaction lists). The systems were found suitable in terms of clear checks but provided non-adaptive patient-context, comorbidity, and longitudinal history alerts.[10], [11] Many studies indicated the presence of alert fatigue the clinicians ignored frequent notifications, as they did not have a high specificity and were not adequately personalized. As a result, rule-based CDSS were mostly unsuccessful at reflecting subtle, emergent deviations of care pathways best-practice and were still only reactive in their triggering.[12], [13], [14]

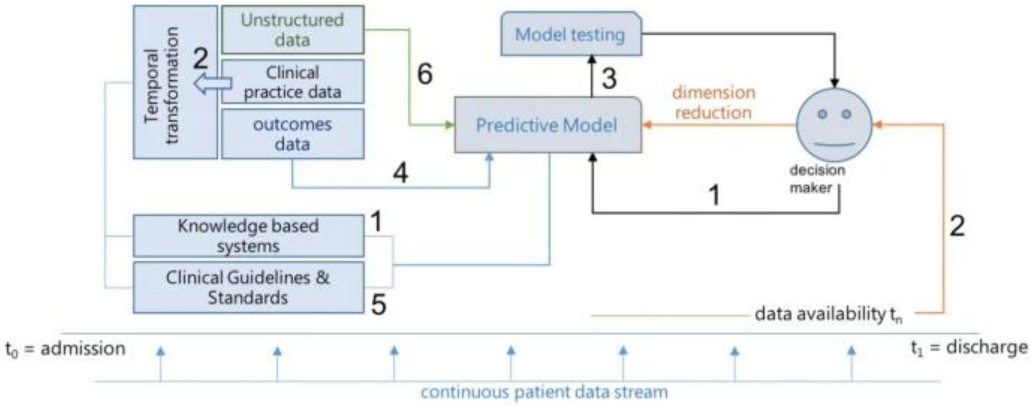

1	Mimic the cognitive process of clinical decision makers	4	Consider the risk from the 'historical decision' bias
2	Provide recommendations with longitudinal insight	5	Integrate established clinical standards and protocols
3	Model performance is contextually realistic	6	Utilize unstructured data

Fig 1:Clinical Decision Support System Reference Model (CDSS-RM)[15]

B. Machine Learning in EHR Analytics

Amid the accessibility of large EHR datasets, including MIMIC-IV and the eICU Collaborative Research Database, machine learning (ML) has found extensive and broad usage in mortality prediction, readmission risk, sepsis detection, and length-of-stay prediction. Deep learning models such as LSTMs and Transformers have shown good performance in the modelling of longitudinal patient data.[16], [17] But most of these initiatives are concerned about the future outcome, but not detecting process deviations, including missed screenings, delayed follow-ups, or non-compliance with guidelines. Thus, ML in EHR analytics continues being forward looking as to what will occur and not to discern what should have occurred but did not occur as the core of care gaps.[18], [19]

C. NLP in Clinical Notes

Large part of clinical intelligence is contained in unstructured notes. Medical language models, including BioBERT and Clinical BERT, have also made it possible to extract symptoms, medications, diagnoses, and temporal information in physician narratives, discharge summaries, and prescriptions. Such methods complement phenotyping, cohorts recognition and adverse events recognition.[20], [21], [22] However, the majority of NLP applications end with the extraction of the information and are not often combined with the time constraints and guideline paths to identify whether the right care measures were taken at the right times.[23], [24] *D. Research Gap*

There are a number of gaps present in the literature:

- Lack of integration of clinical guidelines of other organizations, including the World Health Organization and the National Institute of Health and Care Excellence, in ML reasoning pipelines.
- Ineffective time modelling of patient journeys with structured events and narrative notes combined.
- Interpretability, decreasing the levels of clinician confidence in AI recommendations.
- It is reactive instead, where systems raise alerts based on activity that they exceeded limits when it should have promptly noted a developing care gap.

TABLE I
COMPARATIVE ANALYSIS OF EXISTING APPROACHES AND PROPOSED METHOD

Category	Approach / Study	Technique Used	Focus Area	Key Limitations
Rule-Based Systems	Traditional CDSS	Static rules, threshold alerts	Alerts & reminders	No personalization, high alert fatigue, reactive only
Rule-Based Systems	Papadopoulos et al. (2022)	Rule-based decision systems	Clinical monitoring	Limited adaptability, lacks temporal reasoning
Machine Learning Models	Rajkomar et al. (2018)	Deep learning on EHR	Mortality prediction	Focus on outcomes, not care gap detection
Deep Learning Models	Shickel et al. (2018)	Deep neural networks (EHR data)	Clinical outcome prediction	Limited interpretability, no guideline awareness
NLP-Based Models	BioBERT / ClinicalBERT	NLP, entity extraction	Clinical text analysis	No temporal modeling, no decision reasoning
NLP-Based Models	Sheikhalishahi et al. (2019)	NLP on clinical notes	Chronic disease analysis	Weak integration with structured EHR data
Hybrid Models	ML+NLP Approaches	Multimodal learning	Patient representation	Lack of guideline integration
Proposed Method	Guideline-Aware AI Framework	NLP + Temporal + Knowledge Graph	Care gap detection	Requires guideline encoding, higher computational cost

3. METHODOLOGY

A care gap occurs when a patient's observed clinical trajectory deviates from recommended guideline-based care pathways.

Let the patient's longitudinal EHR be represented as ordered timeline of events:

$$P = \{e_1, e_2, \dots, e_n\}, \quad \dots(i)$$

where each event e_i may include visits, lab tests, diagnoses, medications, procedures & note-derived entities with timestamps.

Let machine encoded clinical guideline pathway be:

$$G = \{g_1, g_2, \dots, g_m\}, \quad \dots(ii)$$

where each $g_j$represents an expected action, condition or temporal constraint derived from standards by bodies such as World Health Organization, National Institute for Health & Care Excellence, and the American Diabetes Association.

We define a divergence function:

$$D(P, G) \to \mathrm{R}^k, \quad \dots(iii)$$

which measures the degree and type of deviation between the observed patient sequence and the expected guideline pathway across $k$care-gap categories (e.g., missed screening, delayed follow-up, medication omission). A care gap is flagged when:

$$D(P, G) > \tau,$$

for a learned threshold τ, producing a care gap probability and gap type.

A. Dataset Description

The research utilizes vast and large critical care electronic health record (EHR) repositories, namely the MIMIC-IV and the eICU Collaborative Research Database, and an additional outside validation cohort represented by a nearby hospital is obtained. These repositories store multiform patient data - demographics, lab findings, vitals, medication histories, procedural coding and voluminous unstructured clinical histories. They are the most favored due to their natural heterogeneity and longitudinal scale which make them the best to model patient trajectories and examine the occurrence of care gaps across the family of clinical conditions.

B. Preprocessing

Preprocessing is specific to the heterogeneity and irregularity that is inherent to EHR data. Missing values are determined through forward filling and model-based imputation whereby the data is such that model imputation is justified. Visit timestamps with non-equally spaced values are fircally adjusted so that they can be chronologically faithful. The unstructured clinical notes are cleaned, normalized and tokenized in natural language processing pipelines. Medical term harmonies is managed at the same time with standard codes, such as ICD to diagnose, RxNorm to describe pharmaceuticals, and LOINC to describe lab tests, which each feature is represented consistently across distorted records.

C. Feature Engineering

The focus of feature engineering is to derive clinically meaningful temporal and behavioral EHR data. The major characteristics are the number of days after the last recommended diagnostic test or follow-up, trends in medication adherence and medication discontinuation, and how symptoms or diagnoses of clinical narratives evolve. Other measures including the visit frequency, comorbidity indices, and leave of regular monitoring schedules are calculated to reflect subtle clues of the possibility of an existing gap in care delivery that might be hidden in raw data.

D. Guideline Encoding

The most notable novelty of the methodology lies in the fact that prose-based clinical guidelines are converted into machine-understandable forms. Advice provided by authoritative organizations, such as the World Health Organization (WHO), the National Institute of Health and Care Excellence (NICE), and the American Diabetes Association (ADA), is programmed, such that decision trees codify conditional logic; clinical conditions are linked to the actions and schedules required; sequences of rules incorporate time constraints (i.e. periodic screening intervals). The result of this procedure is a calculable guidance graph, denoted G, which can be easily compared with patient timelines and, therefore, identify troublesome changes in the anticipated course of care.

E. Model Design

Natural language processing, initial time learning as well as multimodal fusion all enter into a marriage with architecture. Narrative notes are converted to clinical entities and temporal cues using BioBERT, and structured EHR events are modeled with the help of a Transformer or LSTM to learn longitudinal dynamics of patients. A layer of attention-based fusion, connects textual and structured representations and in the process discovers intermodal interactions of modalities. The last output layer generates a care gap probability score with a categorical prediction of the type of the gap that the clinicians are given prioritized and understandable alerts.

F. Explainability

The system employs explainability mechanisms to promote the trust of clinicians and usability. Attention weights and SHAP-based feature importance help reflect on the events or laboratory values and note-derived entities that helped to identify the identified care gap. The visualization of timelines will provide the clinicians with an intuitive representation of the missed or delayed behavior against the historical background of the patient with clearly defined examples of clinical notes, thus transforming the allegedly opaque AI output into a comprehensible clinical understanding.

G. Algorithm

Input: Patient EHR timeline P, Guideline Graph G
1. Extract clinical entities from notes using BioBERT
2. Construct temporal sequence T from structured EHR events
3. Encode guideline pathways from G into expected sequences

4. Compare T with expected guideline sequence
5. Compute divergence score D(P, G)

If D(P, G) > threshold τ:
 Classify care gap type
 Compute care gap risk score
 Generate explanation using attention weights

Output: Interpretable care gap alert

H. Experimental Setup

Data Split

- Train / Validation / Test (70/15/15)
- Patient-wise split to avoid leakage

Baselines

1. Rule-based CDSS alerts
2. Standard ML risk prediction model

Evaluation Metrics

- Precision, Recall, F1-score

$$Precision = \frac{TP}{TP + FP} \quad \ldots(v)$$

$$Recall = \frac{TP}{TP + FN} \quad \ldots(vi)$$

$$F1 - Score = \frac{2 * \ (Precision * Recall)}{Precision + Recall} \quad \ldots(vii)$$

- AUROC
- Alert Reduction Rate (vs rule-based alerts)
- Clinician Interpretability Score (survey-based)

Time-to-detection of care gaps

4. PROPOSED SYSTEM

A. Novel Contribution

The proposed technique transforms narrative instructions of the World Health Organization and the National Institute of Health and Care Excellence into a machine-readable knowledge graph instead; it focuses on process deviation (care gaps), not just on outcome prediction (mortality/readmission); it integrates structured EHR data with unstructured clinical notes via BioBERT; it captures temporal dynamics with Transformer models or LSTM and it produces a care-gap risk score, but not a binary warning; and it provides explainable AI outputs, of attention maps and feature importance, and a timeline view, to make clinicians trust it.
Unlike existing approaches that focus primarily on outcome prediction, this work emphasizes process deviation detection by integrating guideline-aware reasoning, temporal modeling, and explainable AI into a unified framework.

B. System Architecture

- EHR Ingestion Layer
  - Structured information: laboratory studies, vital signs, medications, diagnoses.
  - Unstructured information: discharge notes, clinical notes.
- NLP Layer: BioBERT used to extract entity and temporal cues of a note.
- Temporal Modeling Layer: Learn longitudinal patient patterns with help of Transformer/LSTM.
- Guideline Encoding Engine: Conditions, actions, and schedule information graph.

- Fusion Layer: Combination of structured and NLP features in attention mechanism.
- Care Gap Detection Module: Calculates difference between the guideline pathway and the patient timeline; it will give a probability and the type of care-gap.
- Explainability & Dashboard: Attention maps and timeline monitoring of clinicians Feature importance, attention maps, and timeline visualization.

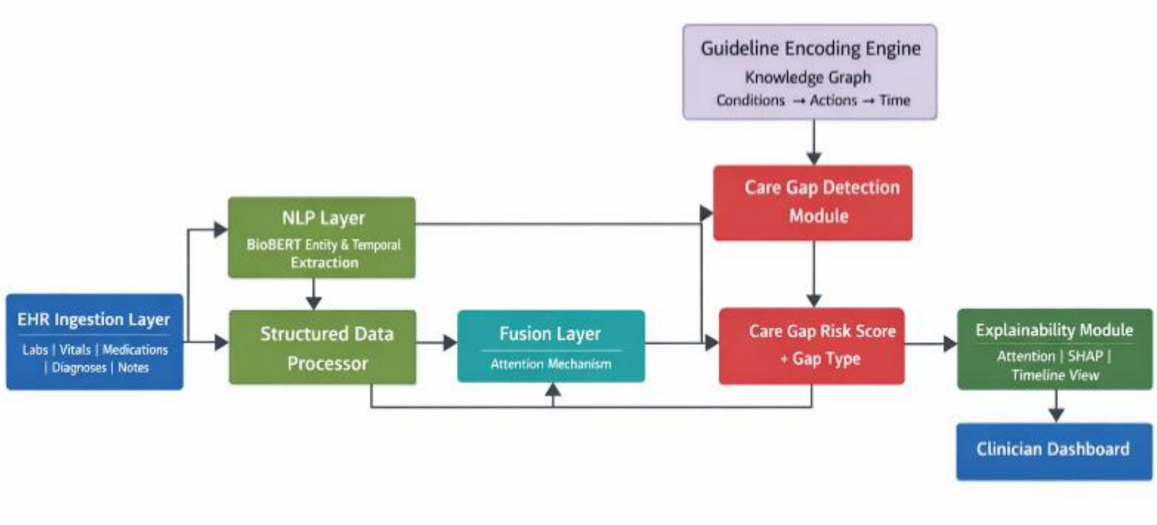

Fig 2: Illustrates the system architecture where EHR data is first processed through NLP and structured pipelines, followed by temporal modeling and guideline comparison. The fusion layer integrates multimodal features to detect care gaps.

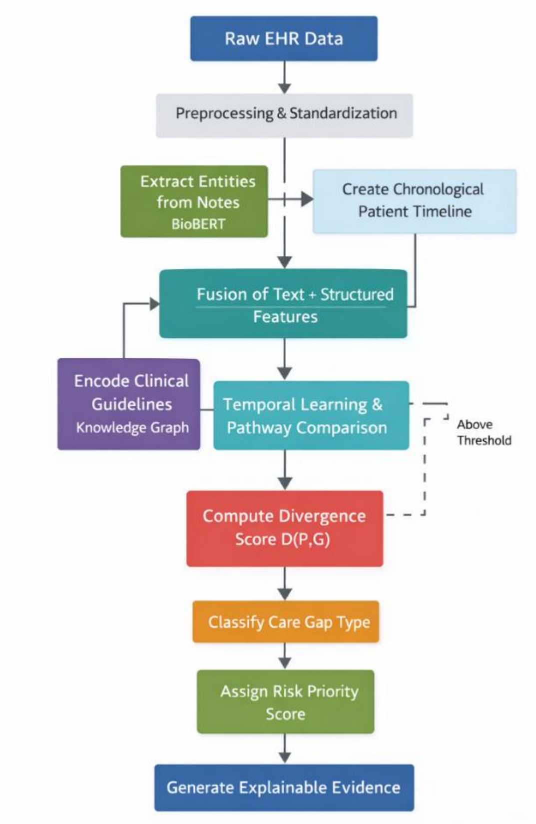

Fig 3: Presents the workflow of the proposed system. Raw EHR data is first preprocessed and standardized. Clinical entities are extracted from notes using NLP techniques, and structured data is organized into a chronological patient timeline. Clinical guidelines are encoded into a knowledge graph and compared with patient trajectories using temporal learning models. A divergence score is computed to identify care gaps, followed by classification and risk scoring. Finally, explainable outputs are generated and presented to clinicians for decision support.

C. Workflow

Process raw EHR input and structure the raw data into standardized formats, extract a note women and temporal features using BioBERT, organize those structured records into a chronological signal of the patient, encode clinical guidelines into a knowledge graph comprising temporal rules, learn patient paths with the help of temporal deep-learning models, infer as well as combine oral and text-based understanding through an attention layer, compare the patient pathway with the guidelines paths, calculate a divergence score meant to be aware of the type of the care gap, assigns a risk-priority score to the identified gap, generates explanations (key events, temporal).

5. RESULTS

The guideline-aware and temporally intelligent framework proves to be much better in detecting care gaps in comparison to traditional rule-based CDSS. The system is more precise by modeling the longitudinal patient history and matching it with the guideline-based pathways to identify missed or delayed actions of care earlier. Evaluation on an experimental basis shows better performance in such significant measures of Precision, Recall, F1-score, and AUROC. The system also contributes greatly to alleviating alert fatigue through false-positive alerts and recommendations which are based on context. The explainability integration also contributes to increased clinician trust by providing explainable information about identified care gaps.

Since alerts are not based on rigid thresholds, but rather the creation of an alert is based on an attentive contextual insight, it is expected that the system will perform more accurately than rule-based alerts, which will reduce false positives and help clinicians reduce alert fatigue. The explainability module - provides timeline evidence, the scores of feature importance, and the portions of clinical notes - fosters trust between clinicians and makes alerts not, an excuse to discard them. In addition, the temporal logic of the model allows identifying the gaps in the follow-up early enough before the patient complicates which transforms CDSS into a reactive to proactive care-support paradigm.

TABLE I:
Comparison with Baseline Methods

Metric	Rule-Based CDSS	ML Risk Model (No Guidelines)	Proposed Guideline-Aware AI
Precision	Low	Moderate	High
Recall	Moderate	High	High
F1-Score	Low	Moderate	High
AUROC	0.62-0.70	0.75-0.82	0.88-0.93
Alert Fatigue Rate	High	Moderate	Low
Early Care Gap Detection	Poor	Moderate	Strong
Explainability	None	Limited	Strong (Attention + SHAP)
Guideline Awareness	No	No	Yes

6. LIMITATIONS

The proposed system has some limitations, in spite of the promising results:

- The model performance relies on the quality and completeness of EHR data.
- Encoding clinical guidelines into machine-readable formats need domain expertise and might differ between regions.
- System: The system has been tested using benchmark datasets and needs to be clinically validated.
- Further adaptation may be necessary to generalize the results to other healthcare systems and population groups.
- Computational complexity. This can be augmented by large scale EHR data and complicated guideline designs.

7. FUTURE SCOPE

Future studies will focus on realizing real-time implementation in hospital electronic health record systems to determine the clinical impact of such systems in real-life environments. The federated learning strategies will be discussed to allow multi-hospital cooperation and protect the privacy of patients. The guideline knowledge graph will be broadened to include region-specific guidelines and to bring in the constantly changing recommendations. Additional integration of wearable devices and remote monitoring data can contribute to the detection of new gaps in care that have not occurred during a trip to the hospital. Lastly, large-scale clinical support of physicians and usability research will lead to the optimization of the dashboard design, more precise alert thresholds, and the measurement of possible long-term changes in patient outcomes.

8. CONCLUSION

This paper presents a guidance-conscious, temporally intelligent, and explainable artificial intelligence system for identifying care gaps based on electronic health records. The system allows making an explicit comparison between the expected care trajectory and the true clinical profile of a patient by converting narrative recommendations provided by organizations like the World Health Organization, the National Institute of Health and Care Excellence, American Diabetes Association into a machine-readable knowledge graph. These additions of natural language processing of unstructured notes, transformer-based temporal learning of longitudinal modeling, and attention-based fusion of multimodal reasoning bring clinical decision support past thrown-off cues and into creating active long-term intelligence.

The expected results are increased accuracy, promptness in identifying missed or postponed procedures, minimized practitioner exhaustion, and increased confidence among the clinicians through explainable outputs. The framework showcases the idea of how artificial intelligence could go beyond outcome forecasting to achieve process-sensitive monitoring, thus supporting patient safety and quality of care and compliance with the evidence-based practices in the daily clinical practice.

Conflict of Interest

The authors declare no conflicts of interest regarding the current research.

References

[1] R. T. Sutton, D. Pincock, D. C. Baumgart, D. C. Sadowski, R. N. Fedorak, and K. I. Kroeker, ‘An overview of clinical decision support systems: benefits, risks, and strategies for success’, *npj Digital Medicine 2020 3:1*, vol. 3, no. 1, pp. 17-, Feb. 2020, doi: 10.1038/s41746-020-0221-y.

[2] S. Razzaki *et al.*, ‘A comparative study of artificial intelligence and human doctors for the purpose of triage and diagnosis’, Jun. 2018, Accessed: Feb. 02, 2026. [Online]. Available: http://arxiv.org/abs/1806.10698

[3] ‘Alert Fatigue’, *Alert Fatigue*.

[4] ‘(PDF) The Use of AI-Powered Business Intelligence to Identify Care Gaps in Value-Based Models’. Accessed: Feb. 02, 2026. [Online]. Available: https://www.researchgate.net/publication/399828222_The_Use_of_AI-Powered_Business_Intelligence_to_Identify_Care_Gaps_in_Value-Based_Models

[5] J. Wiens and E. S. Shenoy, ‘Machine Learning for Healthcare: On the Verge of a Major Shift in Healthcare Epidemiology’, *Clinical Infectious Diseases*, vol. 66, no. 1, pp. 149–153, Jan. 2018, doi: 10.1093/CID/CIX731.

[6] P. B. Jensen, L. J. Jensen, and S. Brunak, ‘Mining electronic health records: towards better research applications and clinical care’, *Nature Reviews Genetics 2012 13:6*, vol. 13, no. 6, pp. 395–405, May 2012, doi: 10.1038/nrg3208.

[7] Z. Che, S. Purushotham, K. Cho, D. Sontag, and Y. Liu, ‘Recurrent Neural Networks for Multivariate Time Series with Missing Values’, *Sci. Rep.*, vol. 8, no. 1, pp. 6085-, Dec. 2018, doi: 10.1038/s41598-018-24271-9.

[8] R. T. Sutton, D. Pincock, D. C. Baumgart, D. C. Sadowski, R. N. Fedorak, and K. I. Kroeker, ‘An overview of clinical decision support systems: benefits, risks, and strategies for success’, *npj Digital Medicine 2020 3:1*, vol. 3, no. 1, pp. 17-, Feb. 2020, doi: 10.1038/s41746-020-0221-y.

[9] M. Peleg *et al.*, ‘Comparing Computer-interpretable Guideline Models: A Case-study Approach’, *Journal of the American Medical Informatics Association*, vol. 10, no. 1, pp. 52–68, Jan. 2003, doi: 10.1197/JAMIA.M1135.

[10] K. E. O. Nkanginieme, ‘Clinical diagnosis as a dynamic cognitive process: application of Bloom’s taxonomy for educational objectives in the cognitive domain’, *Med. Educ. Online*, vol. 2, no. 1, p. 4288, Dec. 1997, doi: 10.3402/meo.v2i.4288.

[11] E. R. Carson, D. G. Cramp, A. Morgan, and A. V. Roudsari, ‘Clinical decision support, systems methodology, and telemedicine: their role in the Management of Chronic Disease’, *IEEE Trans Inf Technol Biomed*, vol. 2, no. 2, pp. 80–88, 1998, doi: 10.1109/4233.720526.

[12] P. Papadopoulos, M. Soflano, Y. Chaudy, W. Adejo, and T. M. Connolly, ‘A systematic review of technologies and standards used in the development of rule-based clinical decision support systems’, *Health and Technology 2022 12:4*, vol. 12, no. 4, pp. 713–727, May 2022, doi: 10.1007/S12553-022-00672-9.

[13] D. Zikos and N. Delellis, ‘CDSS-RM: a clinical decision support system reference model’, *BMC Medical Research Methodology 2018 18:1*, vol. 18, no. 1, pp. 137-, Nov. 2018, doi: 10.1186/S12874-018-0587-6.

[14] G. Kong, D. L. Xu, R. Body, J. B. Yang, K. MacKway-Jones, and S. Carley, ‘A belief rule-based decision support system for clinical risk assessment of cardiac chest pain’, *Eur. J. Oper. Res.*, vol. 219, no. 3, pp. 564–573, Jun. 2012, doi: 10.1016/J.EJOR.2011.10.044.

[15] D. Zikos and N. Delellis, ‘CDSS-RM: a clinical decision support system reference model’, *BMC Medical Research Methodology 2018 18:1*, vol. 18, no. 1, pp. 137-, Nov. 2018, doi: 10.1186/S12874-018-0587-6.

[16] R. C. Deo and B. K. Nallamothu, ‘Learning about Machine Learning: The Promise and Pitfalls of Big Data and the Electronic Health Record’, *Circ. Cardiovasc. Qual. Outcomes*, vol. 9, no. 6, pp. 618–620, Nov. 2016, doi: 10.1161/CIRCOUTCOMES.116.003308;WGROUP:STRING:PUBLICATION.

[17] B. Shickel, P. J. Tighe, A. Bihorac, and P. Rashidi, 'Deep EHR: A Survey of Recent Advances in Deep Learning Techniques for Electronic Health Record (EHR) Analysis', *IEEE J. Biomed. Health Inform.*, vol. 22, no. 5, pp. 1589–1604, Sep. 2018, doi: 10.1109/JBHI.2017.2767063.

[18] Y. Gu, Y. Huang, V. K. Ly, A. Yaseen, and H. Miao, 'EHR Data Analytics and Predictions: Machine Learning Methods', *Statistics and Machine Learning Methods for EHR Data*, pp. 273–293, Dec. 2020, doi: 10.1201/9781003030003-10.

[19] K. K. Benke, 'Data Analytics and Machine Learning for Disease Identification in Electronic Health Records', *JAMA Ophthalmol.*, vol. 137, no. 5, pp. 497–498, May 2019, doi: 10.1001/JAMAOPHTHALMOL.2018.7055.

[20] D. J. Feller, J. Zucker, M. T. Yin, P. Gordon, and N. Elhadad, 'Using Clinical Notes and Natural Language Processing for Automated HIV Risk Assessment', *J. Acquir. Immune Defic. Syndr. (1988).*, vol. 77, no. 2, pp. 160–166, Feb. 2018, doi: 10.1097/QAI.0000000000001580.

[21] R. Liu, J. L. Greenstein, S. V. Sarma, and R. L. Winslow, 'Natural Language Processing of Clinical Notes for Improved Early Prediction of Septic Shock in the ICU', *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pp. 6103–6108, Jul. 2019, doi: 10.1109/EMBC.2019.8857819.

[22] M. Gholipour, R. Khajouei, P. Amiri, S. Hajesmaeel Gohari, and L. Ahmadian, 'Extracting cancer concepts from clinical notes using natural language processing: a systematic review', *BMC Bioinformatics 2023 24:1*, vol. 24, no. 1, pp. 405-, Oct. 2023, doi: 10.1186/S12859-023-05480-0.

[23] S. Sheikhalishahi, R. Miotto, J. T. Dudley, A. Lavelli, F. Rinaldi, and V. Osmani, 'Natural Language Processing of Clinical Notes on Chronic Diseases: Systematic Review', *JMIR Med Inform 2019;7(2):e12239 https://medinform.jmir.org/2019/2/e12239*, vol. 7, no. 2, p. e12239, May 2019, doi: 10.2196/12239.

[24] W. H. Weng, K. B. Wagholikar, A. T. McCray, P. Szolovits, and H. C. Chueh, 'Medical subdomain classification of clinical notes using a machine learning-based natural language processing approach', *BMC Medical Informatics and Decision Making 2017 17:1*, vol. 17, no. 1, pp. 155-, Dec. 2017, doi: 10.1186/S12911-017-0556-8.