Vol. 8, No. 2, April 2026

ISI Journal of Information Systems and Informatics

Published By Asosiasi Doktor Sistem Informasi Indonesia

# Sentiment Analysis in Electronic Health Records for Patient-Centric Care: A Systematic Literature Review of Methods, Applications, and Challenges

**Caroline Mhlanga[1], Belinda Ndlovu[2]**

[1,2] Department of Informatics and Analytics, National University of Science and Technology, Bulawayo, Zimbabwe

**Received:**
December 2, 2025
**Revised:**
March 10, 2026
**Accepted:**
April 1, 2026
**Published:**
April 12, 2026

Corresponding Author:
**Author Name*:**
Belinda Ndlovu
**Email*:**
belinda.ndlovu@nust.ac.zw

DOI:
10.63158/journalisi.v8i2.1545

© 2026 Journal of Information Systems and Informatics. This open-access article is distributed under a (CC-BY License)

**Abstract.** This study examines the role of sentiment analysis in EHR narratives for enhancing patient-centred care, focusing on methodological approaches, application domains, and implementation challenges in clinical settings. A systematic literature review (SLR) was conducted in accordance with PRISMA guidelines. Relevant studies were retrieved from Scopus, Web of Science, IEEE Xplore, and PubMed. The search, conducted in September 2025, included peer-reviewed articles published between 2021 and September 2025. The findings reveal a clear shift from rule-based and traditional machine learning approaches to transformer-based models. Sentiment analysis is increasingly applied in areas such as mental health, oncology, and patient experience monitoring. However, most implementations remain domain-specific and are not fully integrated into routine clinical workflows. This study provides a structured synthesis of sentiment analysis in EHRs and identifies key gaps between methodological advancements and real-world implementation. It advances a socio-technical perspective that integrates analytical performance, clinical applicability, and governance considerations, offering a consolidated lens for understanding sentiment-aware healthcare systems. Despite rapid methodological progress, the impact of sentiment analysis in EHRs remains constrained by limited scalability and insufficient integration into clinical practice.

**Keywords:** ClinicalBERT, Sentiment Analysis, Electronic Health Records (EHR), Clinical Natural Language Processing (Clinical NLP), Transformer Models

## 1. INTRODUCTION

Electronic Health Records (EHRs) have been at the centre of modern healthcare information infrastructure, capable of recording systematic health histories, diagnoses, treatments, and outcomes [1], [2], [3]. EHRs also include unstructured biomedical data, such as rich narrative text in the form of progress notes, nursing observations, discharge summaries, and patient comments on different aspects of their lived experience, concerns, and emotional conditions [1]. The affective content of EHRs, such as anxiety, fear, frustration, hopelessness, etc., is now increasingly recognized as an important factor influencing health outcomes and patient compliance with treatment regimens [4], [5]. The requirement for a systematic interpretation of latent emotional cues embedded in EHRs has now emerged as a significant healthcare informatics issue with the shift towards patient-centered care models [6], [7], [8].

Sentiment analysis, a natural language processing (NLP) task, can provide computational tools for extracting subjective content, such as emotions, opinions, and sentiments, from unstructured text data [6], [7]. When applied to EHRs, sentiment analysis enables integration of psychosocial data into clinical workflows, identifying psychological and emotional patterns and complementing structured clinical data with experience data from unstructured EHRs [9], [10]. Breakthroughs in machine learning techniques, especially the emergence of transformer models such as ClinicalBERT, a domain-specific architecture, will further accelerate understanding of EHRs by enabling deeper analysis of clinical narratives [11], [12]. Despite technological advancements, a gap remains in the integration of sentiment analytics into EHRs.

Traditional sentiment analysis processing pipelines have numerous limitations in the healthcare context [2], [8], [13]. Traditional or lexicon-driven methods struggle to process medical jargon, negative statements, and implicit emotional expressions. At the same time, modern deep learning and transformer models have raised numerous issues regarding requirements, complexity, and interpretability, as well as dataset bias [14], [15]. In addition, most of the literature reviewed emphasizes the algorithm and the alignment of sentiment analysis results with patient-centred goals and ethical rules for processing emotive, sensitive information [16], [17], [18].

Despite growing interest in clinical NLP, existing studies [3], [7], [15] are scattered and varied in terms of methods and domains. This study fills this void by offering a synthesized view that combines model performance and clinical applicability. As a result, this study presents a synthesized view of existing capabilities and future prospects for clinical informatics and digital health practitioners.

This review will proceed with the following research questions as follow;

1. What models of NLP and frameworks have been used in the extraction and analysis of patient sentiment from Electronic Health Records?
2. How have application domains and use cases from sentiments enhanced patient-centric results in a healthcare context?
3. What opportunities exist for leveraging sentiment analysis to enhance patient-centered care within EHR ecosystems?
4. What challenges exist for the reliability, generalisability, or applicability of sentiment models for EHRs?

This study makes three key contributions. First, it provides a focused synthesis of sentiment analysis approaches applied to electronic health record narratives within the context of patient-centred care. Second, it identifies key methodological, contextual, and governance-related gaps that limit real-world clinical integration. Third, it develops a socio-technical conceptual framework grounded in the reviewed literature to guide the responsible and scalable implementation of sentiment-aware healthcare systems.

## 2. METHODS

The study employed the SLR methodology, guided by the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 framework. The study adopted the PRISMA framework because of its ability to guide the inclusion of empirical studies systematically [15], [19], [20].

### 2.1. Systematic Review Workflow

The review followed a structured multi-stage workflow comprising (1) database search execution, (2) duplicate removal using Mendeley reference management software, (3) title and abstract screening based on predefined inclusion criteria, (4) full-text eligibility

assessment, (5) data extraction using a structured template, (6) quality appraisal, and (7) thematic synthesis. Each stage was conducted systematically to ensure transparency and reproducibility of the review process. This structured process ensures transparency, reproducibility, and methodological rigour in line with established systematic review guidelines [19], [21], [22]. Figure 1 shows the PRISMA flow diagram for this SLR.

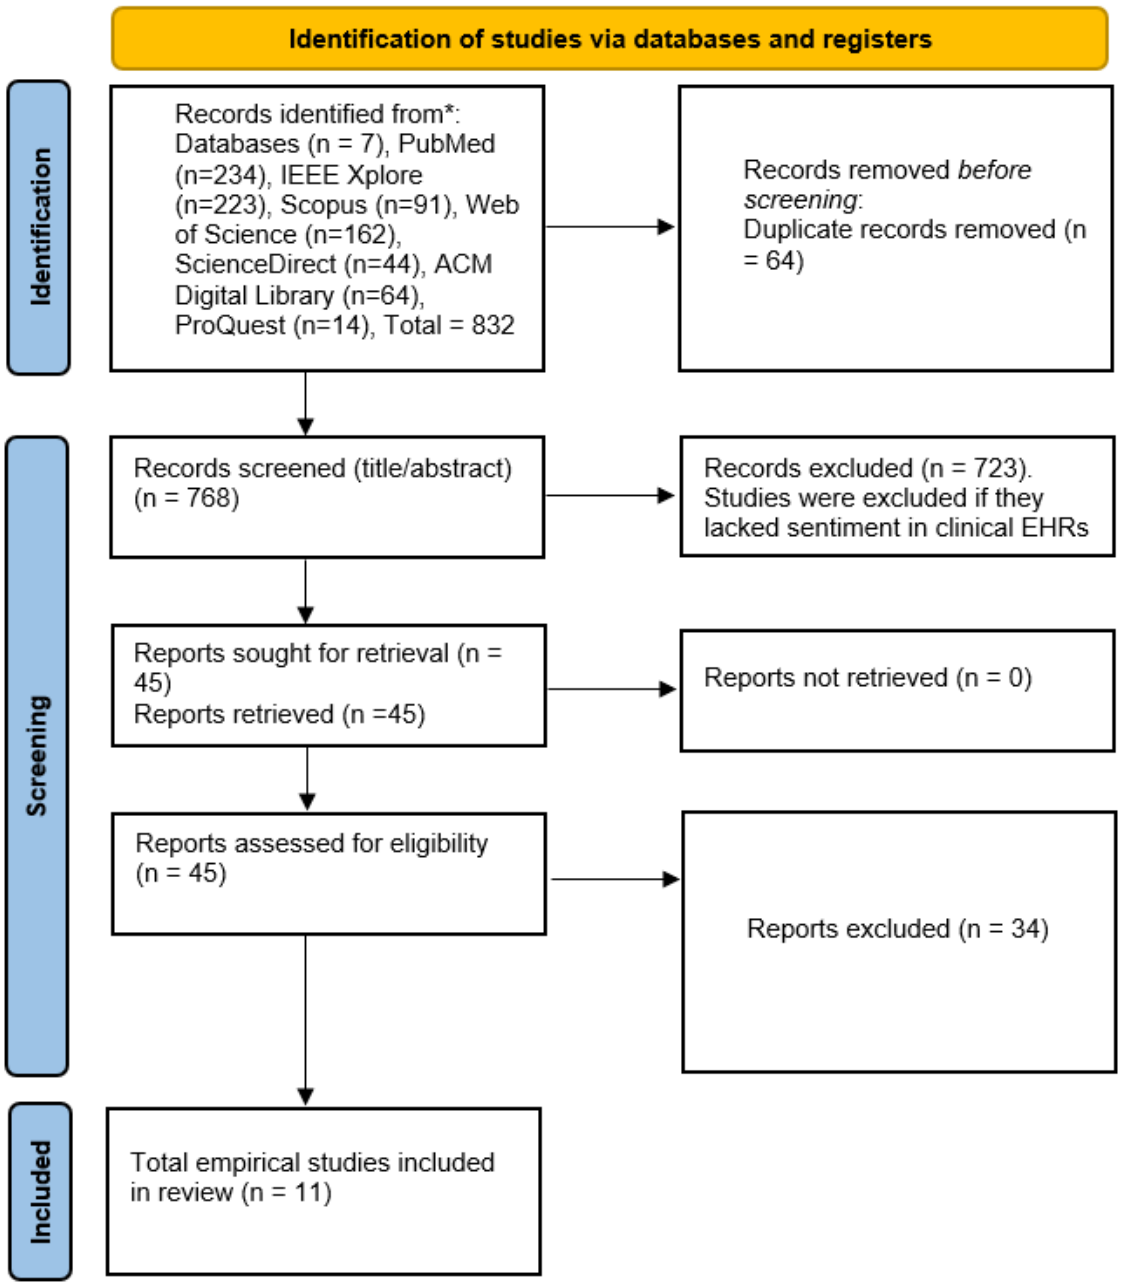



**Figure 1.** PRISMA Flow Diagram for Study Selection

### 2.2. Database Search Strategy

A comprehensive, systematic database search was conducted in September 2025 to identify peer-reviewed studies on AI-based sentiment analysis in EHRs. Seven major scholarly databases were selected for their multidisciplinary scope and for indexing both clinical and computational literature, including PubMed, IEEE Xplore, Scopus, Web of Science, ScienceDirect, ACM Digital Library, and ProQuest. The following Boolean search expression was used consistently across all databases, with adjustments for syntax variations:

*("sentiment analysis" OR "opinion mining" OR "emotion detection" OR "affective computing" OR "natural language processing" OR "text mining") AND ("electronic health record*" OR "EHR" OR "clinical notes" OR "medical records" OR "health data" OR "patient records") AND ("patient-centric care" OR "patient-centred care" OR "personalized medicine" OR "patient engagement" OR "empathy" OR "patient experience").*

The search yielded a total of 832 records: 234 from PubMed, 223 from IEEE Xplore, 91 from Scopus, 162 from Web of Science, 44 from ScienceDirect, 64 from the ACM Digital Library, and 14 from ProQuest. The search was restricted to peer-reviewed journals and conference papers from 2021 to 2025 to focus on the latest advancements in NLP and AI-based sentiment analysis. Studies were limited to English for interpretive consistency and clarity.

### 2.3. Inclusion and Exclusion Criteria

The inclusion and exclusion criteria for this SLR were to ensure conceptual and methodological consistency. The inclusion criteria consisted of empirical studies published in peer-reviewed journals between 2021 and 2025 and written exclusively in English. For this purpose, eligible studies included only published journal articles and conference papers that applied sentiment analysis, emotion detection, affective computing, or NLP to EHRs or closely related clinical text sources. To align with patient-centred healthcare objectives, studies needed to explicitly explore outcomes related to patient experience, emotional well-being, empathy, personalized care, or sentiment-informed clinical decision support within clinical information systems.

Conversely, non-peer-reviewed publications, preprints, editorials, commentaries, conceptual papers, and grey literature were excluded. Only studies that do not incorporate emotion- or sentiment-based analysis in their disease prediction or risk modelling were excluded. Moreover, studies informed solely by non-clinical data sources, such as social media platforms or online forums, have been excluded. Similarly, to avoid overlap in synthesis, secondary studies comprising systematic reviews and meta-analyses are excluded, as are publications that are not available in full text or are published in a language other than English.

### 2.4. Screening

The selection process adhered to the PRISMA guidelines described by [23] to ensure a feasible, highly relevant pool of eligible studies. The primary search produced 832 records. After removing 64 duplicate records in Mendeley, the pool comprised 768 unique records. To prevent an impractical number of full-text screenings, a rigorous title and abstract screening was conducted. Records were removed if they failed to focus on the sentiment component in a clinical EHR setting. This step led to the removal of 723 records. The remaining 45 records were screened thoroughly. After the full-text screening, 34 records were removed for lacking empirical validation of sentiment, being non-EHR datasets, or being purely conceptual. This resulted in the synthesis of 11 primary studies. The peer-reviewed empirical criterion was strictly applied, excluding any papers that were purely conceptual or lacked primary data validation.

### 2.5. Quality Assessment

All included studies (n = 11) were appraised for quality using predefined criteria, ensuring consistency and methodological transparency across the review. This was done by following the recommended approaches for conducting systematic literature reviews in the fields of information systems and clinical informatics. The quality of each study was evaluated according to the following criteria: (1) methodological clarity, (2) transparency of data, (3) model evaluation, and (4) reproducibility. The quality of the studies was not a factor in exclusion; instead, it was considered during interpretation of the results. Overall, the quality of the included studies was moderate, with sufficient model evaluation but limited transparency in data sourcing and reproducibility. This is reflected in the study's results. The results of the quality assessment of the studies were considered during the

thematic synthesis and the development of the proposed conceptual framework. A total of 11 papers met the final inclusion criteria. These studies are presented in Table 1.

**Table 1.** Quality Appraisal of Included Studies

| Author(s) | Methodological Clarity | Data Transparency | Model Evaluation | Reproducibility |
|---|---|---|---|---|
| [24] | High | Medium | High | Medium |
| [25] | High | Low | High | Low |
| [26] | Medium | Medium | Medium | Medium |
| [9] | High | Medium | High | Medium |
| [27] | High | High | Medium | High |
| [28] | High | Medium | High | Medium |
| [2] | Medium | Medium | Medium | Low |
| [29] | High | Low | High | Low |
| [30] | Medium | Medium | Medium | Medium |
| [14] | High | High | Medium | High |
| [31] | High | Medium | High | Medium |

## 3. RESULTS AND DISCUSSION

### 3.1. Data Reporting

An eight-column table was designed to extract relevant information for the SLR. A total of 11 publications were analyzed in this study, as shown in Table 2.

**Table 2.** Papers that met the inclusion criteria

| Ref | Country | Models of NLP | Frameworks | Domain | Algorithms / Techniques | Opportunities | Challenges |
|---|---|---|---|---|---|---|---|
| [23] | USA | • ClinicalBERT (Transformer-based). | Predictive Analytics | Outcome Prediction | • Bag-of-Words (BoW)<br>• ClinicalBERT<br>• MTTL-ClinicalBERT | Clinical Decision Support | Linguistic variation |
| [24] | USA | MTTL-ClinicalBERT | Affective Computing | Clinical Classification | • Bi-LSTM | Patient-led monitoring | Data scarcity |

| Ref | Country | Models of NLP | Frameworks | Domain | Algorithms / Techniques | Opportunities | Challenges |
|---|---|---|---|---|---|---|---|
| | | | | | • Attention Mechanism | | |
| [25] | UK | Descriptive NLP. | Clinical Decision Support | Care Optimization | • BERT<br>• Random Forest<br>• SVM | Real-time risk alerts | Model transparency |
| [9] | USA | • BERT<br>• spaCy. | Service Quality | Clinical Decision Support | • VADER<br>• TextBlob | Service optimization | Lexicon limitations |
| [26] | India | BERTopic | • VADER<br>• NRCLex | Service Evaluation | • RoBERTa<br>• Zero-shot Learning | Goal-of-care accuracy | Semantic complexity |
| [27] | USA | Deep-learning NLP. | Serious Illness Framework | Palliative Care | • CRFs<br>• LDA | Outcome prediction | Entity ambiguity |
| [2] | USA | Descriptive NLP annotation | Patient-Centered Care (PCC) | Patient Experience | • VADER<br>• NRC Lexicon | Telehealth refinement | UX data noise |
| [28] | USA | Open-source NLP. | VADER Sentiment Analyzer. | Quality Improvement | • Latent Dirichlet Allocation (LDA)<br>• BERT<br>• BioBERT | Precision phenotyping | Narrative sparsity |
| [29] | Netherlands | Text mining algorithms | Symptom-aware Mining | Clinical Diagnostics | • VADER<br>• SentiWordNet | Satisfaction insights | Clinical jargon |
| [14] | UK | GATE (General Architecture for Text Engineering) | Social Context Mining | Mental Health | • Pattern Matching<br>• NegEx | Phenotype discovery | Rigid rule-sets |
| [30] | USA | ML-based NLP pipeline | Personalized Medicine Focus | Oncology | • SVM<br>• Random Forest | Psychosocial mapping | Contextual nuance |

| Ref | Country | Models of NLP | Frameworks | Domain | Algorithms / Techniques | Opportunities | Challenges |
|---|---|---|---|---|---|---|---|
| | | | | | • doc2vec<br>• fasttext<br>• XGBoost<br>• CNN | | |

### 3.2. Publication Trends

Figure 2 presents the publication trend of the studies reviewed in this SLR from 2021 to 2025. From the distribution, it can be observed that the research output was moderate in 2021 when n=3. This was reduced in 2022 to n=1, and then increased in 2023 to n=2. This was followed by a peak in research output in 2024 (n=2) and in 2025, which could indicate the peak research output within the reviewed period, possibly due to a sudden surge in interest in patient-centric natural language processing and sentiment-aware analytics of electronic health records. From the quantitative evidence, it can be observed that there was a definitive shift from exploratory research in the early days of documentation to system-level research in virtual healthcare, mental health analytics, knowledge representation, and ethical AI.

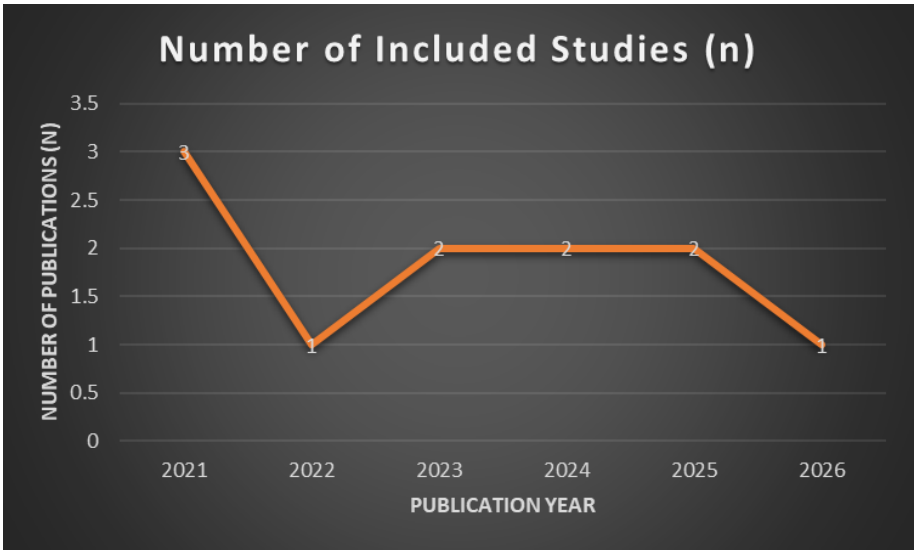



Figure 2: Publication trend

The most trending topic during this analytical period is the shift from basic sentiment polarity to high-density predictive signalling using transformer-based attention

mechanisms, such as "ClinicalBERT," to tap into "dark data" such as family discussions (obtaining a performance boost to an AUC of 0.92). The number of studies published after 2023 indicates that sentiment analysis of EHRs is shifting into a full-blown field of study, often envisioned as a requirement for empathetic and governance-aware digital health. The publication trend shows a gradual increase from 2021 to 2025, with a modest peak in later years. However, given the small number of studies included (n=11), these trends should be interpreted cautiously.

### 3.3. Study characteristics and origin

Figure 3 provides a quantitative overview of the structural characteristics and geographical scope of the 11 major empirical research studies, evaluating the institutional and regional landscapes that currently frame sentiment-aware Electronic Health Record research.

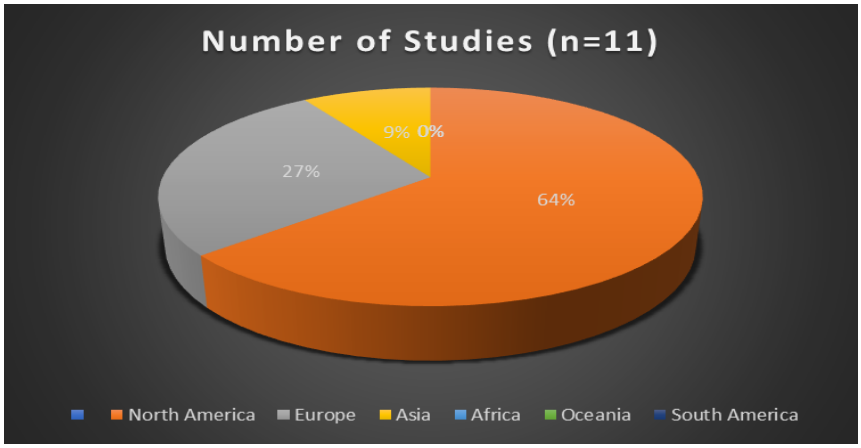



Figure 3: Study characteristics and origin

Quantitative synthesis of the 11 primary studies revealed a heavy research concentration in the Global North, with North America dominant at 63.6% (n=7), followed by Europe at 27.3% (n=3). The United States clearly outperforms other countries with n=7, leveraging existing data infrastructures such as MIMIC-III to prioritize outcome prediction and palliative care. Meanwhile, European countries, such as the UK, have shown a clear preference for mental health analytics and social determinants extracted from narratives. Asia, at 9.1% (n=1), continues to prioritize EHR service evaluation and user experience. One of the most notable findings is the "digital divide" between nations, with no representation from Africa, South America, or Oceania. It can be assumed that this is

due to the existing infrastructure for narrative-rich EHR systems and NLP models that support diverse regional dialects. The implications for big data science are clear: although sentiment analysis is now an advanced field in Western clinical settings, it is not yet globally applicable. To create more patient-centric care, there is an imperative for cross-lingual and low-resource frameworks to incorporate these "missing" perspectives from the Global South to inform digital health governance.

### 3.4. Models of NLP

Table 3 shows the distribution of NLP models across the 11 studies and reveals a sophisticated transition from traditional linguistic theory to high-density deep learning architectures. To ensure transparency in the methodology used, these models are organized into main NLP Model categories and their specific Sub-Model implementations, as identified in Table 1.

**Table 3:** Models of NLP

| Main Models of NLP | Sub-Models of NLP | Studies |
|---|---|---|
| Transformer-based Architectures | • ClinicalBERT (Transformer-based)<br>• BERT<br>• spaCy<br>• BERTopic<br>• MTTL-ClinicalBERT | [9], [23], [24], [26] |
| Rule-based and Lexicon Models | • Descriptive NLP annotation<br>• Text mining algorithms<br>• GATE (General Architecture for Text Engineering) | [2], [14], [29] |
| Hybrid and Statistical ML Pipelines | • Descriptive NLP<br>• Deep-learning NLP<br>• ML-based NLP pipeline<br>• Open-source NLP | [25], [27], [28], [30] |

The quantitative analysis also reveals that the most prominent paradigm is Transformer-based Architectures, which are employed in 36.4% (n=4) of the corpus, owing to their superior ability to perform context-aware bidirectional encoding, which is critical for

decoding the dark data associated with complex clinical narratives. Moreover, 27.3% (n=3) of the studies employed combined theories, such as the hybrid model of BERTopic and VADER, to combine the power of statistical topic modelling and affective polarity [26]. The most glaring research gap in the existing corpus is the explainability of the most widely used Transformer-based Architectures, such as ClinicalBERT, which, although it attains high AUC values, remains a black box in its decision-making process. Therefore, it is recommended that the most critical aspect for future research is the integration of Explainable AI with Transformer-based Architectures to ensure that sentiment-driven insights are clinically actionable and transparent to digital health governance.

### 3.5. Frameworks

As indicated (Table 4), the frameworks were grouped into main frameworks and their corresponding sub-frameworks.

Table 4: Frameworks

| Main Frameworks | Sub-Frameworks | Studies |
|---|---|---|
| Clinical Informatics and Decision Support | • Predictive Analytics<br>• Clinical Decision Support<br>• VADER<br>• NRClex<br>• Symptom-aware Mining | [23], [25], [26], [29] |
| Patient-Centric and Experience Mining | • Service Quality<br>• Patient-Centered Care (PCC)<br>• Personalized Medicine Focus | [2], [9], [30] |
| Affective and Social Contextualization | • Affective Computing<br>• Serious Illness Framework<br>• Social Context Mining | [14], [24], [27] |

The quantitative synthesis reveals the two-way dominance of Clinical Informatics and Patient-Centric paradigms, each accounting for 36.4% (n=4) of the corpus. This trend indicates that the current state of Big Data Science in EHR research is shifting away from isolated polarity metrics towards models that prioritize clinical relevance and the need for personalized healthcare delivery. Notably, 18.2% (n=2) of the studies employed a combined-theory approach, integrating Service Evaluation models with the VADER and

NRCLex sub-models to bridge the gap between technical sentiment analysis and experiential service data [27]. These findings point to a significant research gap: while Affective Computing and Predictive Analytics models are highly effective at detecting emotional signals, there is a clear lack of Sentiment Governance Frameworks. This suggests that while current models are highly effective in detecting empathy or distress, there is no concrete framework to properly govern the ethical and legal integration of these dark data points into the official medical record. As such, future studies must advance from detection to the development of hybrid models that integrate Transformer-based sentiment analysis with formal digital health governance to ensure proper data integrity and clinical actionability.

### 3.6. Application Domain and Use Case

The thematic environment of the 11 included studies (Table 5), reflects a deliberate shift from administrative auditing to clinical utility. For methodological transparency, these are split into main and sub-application domains and use cases.

**Table 5.** Application Domain and Use Case

| Main Application Domain and Use Cases | Sub-Application Domain and Use Case | Studies |
|---|---|---|
| Clinical Diagnostics and Care Optimization | • Outcome Prediction<br>• Clinical Classification<br>• Care Optimization<br>• Clinical Diagnostics<br>• Oncology | [23], [24], [25], [29], [30] |
| Patient Experience and Service Evaluation | • Clinical Decision Support<br>• Service Evaluation<br>• Patient Experience<br>• Quality improvement | [2], [9], [26], [28] |
| Specialized Mental and Palliative Care | • Palliative Care<br>• Mental Health | [14], [27] |

The results of the quantitative analysis show that the most prevalent categories are Clinical Diagnostics and Care Optimization, accounting for 45.4% (n=5) of the data. The prevalence of these categories suggests that the analysis of sentiment is becoming increasingly prevalent for sophisticated risk assessment, such as mortality prediction and oncological treatment monitoring [14], [23]. Moreover, 18.2% (n=2) of the studies showed combined theoretical applications, combining Service Evaluation with Clinical Decision Support to close the gap between technical data and qualitative patient feedback [26]. There is a clear research gap with respect to the use of these models in Primary Care and Preventive Medicine, with the current literature being predominantly focused on reactive, acute, or chronic disease states. This is an indication that the use of sentiment-aware EHRs for the sub-clinical prediction of patient disengagement or psychological deterioration has yet to be explored. Therefore, it is recommended that future studies extend into longitudinal primary care data to enable a shift from reactive modelling to proactive, patient-focused interventions in digital health governance.

### 3.7. Algorithms / Techniques of machine learning

Table 6 shows the distribution of algorithms and techniques in the 11 studies included, which indicates a shift in strategy from keyword-based mining to more complex attention mechanisms. These are grouped under the main algorithms/techniques and their corresponding sub-algorithms/techniques.

**Table 6.** Algorithms / Techniques

| Main Algorithms / Techniques | Sub-Algorithms/ Techniques | Studies |
|---|---|---|
| Lexicon and Rule-Based Implementation | • VADER<br>• TextBlob<br>• NRC Lexicon<br>• SentiWordNet<br>• Pattern Matching<br>• NegEx | [2], [9], [14], [29] |
| Transformer and Attention Architectures | • Bag-of-Words (BoW)<br>• ClinicalBERT<br>• MTTL-ClinicalBERT<br>• Bi-LSTM<br>• Attention Mechanism | [23], [24], [25], [26], [28] |

| Main Algorithms / Techniques | Sub-Algorithms/ Techniques | Studies |
|---|---|---|
| | • BERT<br>• RoBERTa<br>• Zero-shot Learning<br>• BioBERT | |
| Statistical ML and Dimensionality Reduction | • Random Forest<br>• SVM<br>• CRFs<br>• LDA<br>• doc2vec<br>• fasttext<br>• XGBoost<br>• CNN | [25], [27], [28], [30] |

The quantitative results show a two-way dominance of Lexicon-based and Transformer-based approaches, with each type found in approximately 45.4% (n=5) of the literature. Although VADER and SentiWordNet lexicons are necessary for the transparency required in service assessment, attention-based models and ClinicalBERT have set the standards for the performance of mortality and oncology risk assessment [23], [24]. Interestingly, 18.2% (n=2) of the literature used a combination of theories, specifically LDA (Latent Dirichlet Allocation) with transformer-based embeddings, to integrate thematic complexity with affective polarity. There is a vast research gap in Dynamic Sentiment Calibration; the current state of the art is grossly cross-sectional, with no attempt to capture the dynamic instability of patient sentiment throughout the care continuum. It is proposed that future research focus on recurrent attention models and data pipelines to improve the temporal resolution of analysis, thus promoting a more dynamic framework for digital health governance.

### 3.8. Opportunities

As depicted in Table 7, the thematic analysis of opportunities identified in the 11 included studies indicates a significant trend towards integrating deep learning insights into clinical processes to improve patient outcomes. The opportunities are major and sub-opportunities.

**Table 7.** Opportunities

| Main Opportunities | Sub-Opportunities | Studies |
|---|---|---|
| Clinical Decision Support (CDS) and Risk Management | • Clinical Decision Support<br>• Patient-led monitoring<br>• Real-time risk alerts<br>• Service optimization | [9], [23], [24], [25] |
| Care Optimization and Advanced Phenotyping | • Goal-of-care accuracy<br>• Outcome prediction<br>• Precision phenotyping<br>• Phenotype discovery | [14], [26], [27], [28] |
| Service Refinement and Psychosocial Mapping | • Telehealth refinement<br>• Satisfaction insights<br>• Psychosocial mapping | [2], [29], [30] |

The results of the quantitative synthesis show a two-fold dominance of clinical decision support and care optimization, which together comprise 36.4% (n=4) of the corpus. This is because of the pressing clinical need to leverage affective signals, such as patient-reported distress or goal incongruence, as high-density predictive features for diagnostic validity and mortality risk [9], [23]. Interestingly, 27.3% (n=3) of the literature adopted a hybridized strategy that combined NLP with patient-centred care and social context theories to improve the scalability of psychosocial mapping and telehealth service audits. However, a substantial research gap remains in real-time closed-loop feedback systems. Although existing models have identified the potential for real-time alerts, there is a conspicuous absence of evidence regarding their implementation in active bedside settings. Therefore, future studies should focus on designing and developing Interactive Affective Dashboards that integrate transformer-based sentiment analysis outputs into EHR systems to promote empathetic and data-driven digital health stewardship.

### 3.9. Challenges

The technical and operational challenges identified across the 11 studies in Table 8 concern the processing of unstructured clinical narratives into big data. The challenges are classified into Main Challenges and their respective Sub-Challenges.

**Table 8.** Challenges

| Main Challenges | Sub-Challenges | Studies |
|---|---|---|
| Linguistic and Semantic Complexity | • Linguistic variation<br>• Semantic complexity<br>• Entity ambiguity<br>• Clinical jargon<br>• Contextual nuance | [23], [26], [27], [29], [30] |
| Data Governance and Technical Constraints | • Data scarcity<br>• Model transparency<br>• Lexicon limitations<br>• Narrative sparsity<br>• Rigid rule-sets | [9], [14], [24], [25], [28] |
| Operational Noise | • UX data noise | [2] |

Quantitative meta-analysis reveals a two-way dominance of linguistic complexity and data governance limitations, each contributing 45.4% (n=5) to the list of challenges. This symmetry highlights a pressing need for big data science: although Transformer models address entity vagueness, they also worsen the problem of model interpretability, thereby creating a trade-off between predictive accuracy and clinical validity [6], [9]. Moreover, 18.2% (n=2) of the text corpus exhibited overlapping challenges related to linguistic diversity and narrative density, especially in technical domains such as oncology and chronic illnesses. These findings demonstrate a fundamental research gap: a lack of analysis within the NLP-EHR framework, despite acknowledgement of semantic complexity. This lack of analysis within the framework suggests that the clinical impact of misclassification remains unknown. As a result, there is a need to develop self-correcting algorithms and uncertainty-driven deep learning frameworks to ensure the reliability, explainability, and achievability of sentiment-driven knowledge in digital health governance.

Holistically, a comparative analysis across the included studies reveals a strong dominance of transformer-based models, particularly ClinicalBERT, which can be attributed to their superior ability to capture contextual nuances in clinical narratives. However, this dominance also reflects a broader methodological convergence in the field,

potentially limiting exploration of alternative approaches. Additionally, the concentration of studies in specific domains such as oncology and palliative care suggests that sentiment analysis applications are currently driven by areas where patient experience is most critical, rather than being evenly distributed across healthcare contexts.

### 3.10. Discussion

This section interprets the review's findings in relation to the study objectives, highlighting key methodological trends, application domains, and implementation challenges. The discussion integrates insights across the reviewed studies to provide a coherent understanding of how sentiment analysis is evolving within EHR systems and its implications for patient-centered care.

#### 1) Models of NLP used in the extraction and analysis of patient sentiment from Electronic Health Records

The results from the analysis of the studies reviewed above clearly indicate a shift in the methods used to extract and analyze patient sentiment from Electronic Health Records (EHRs), from rule-based approaches to more sophisticated transformer-based models in patient-centric analytical environments. In the studies reviewed above, three major categories of NLP models were identified: (1) Transformer models, (2) Rule-based models, and (3) Hybrid models. These models are embedded in more comprehensive computational and organizational structures intended to facilitate patient-centric analytics [9], [11], [31], [32].

**Transformer models**: The synthesis indicates that transformer-based architectures have become the dominant approach for sentiment extraction in EHR narratives, with domain-adapted BERT variants, multi-task learning models, and contextual embedding techniques demonstrating superior performance [3], [6]. In particular, adaptations such as ClinicalBERT and specialized transformer models have shown enhanced capability in handling negation, temporal expressions, and implicit affect embedded in clinical text [24], [30]. Evidence from domains such as mental health and goals-of-care documentation further suggests that attention mechanisms improve contextual disambiguation, leading to higher recall and F1 scores in sentiment classification tasks [25], [27]. This is especially relevant in healthcare, where meaning is often nuanced and context-dependent. However, while transformer-based models offer improved scalability and generalisability

compared to traditional approaches [9], [33], their deployment remains largely experimental. This highlights a persistent gap between methodological advancement and clinical implementation, reinforcing the need to align model development with real-world healthcare constraints.

**Rule-based Models**: The findings reaffirm the foundational role of rule-based and lexicon-driven models in the early development of sentiment analysis for EHR systems. These approaches, including ontology-aligned keyword extraction, domain-specific polarity dictionaries, and pattern recognition techniques, provided structured, interpretable mechanisms for analyzing patient feedback and mental health narratives [28], [32]. Their strength lies in their precision and transparency, which are particularly valuable in clinical contexts where explainability is essential. However, the synthesis indicates that these models are inherently limited in handling contextual complexity, especially in cases involving ambiguous language, physician hedging, and temporality within clinical documentation [2], [13], [14]. Despite these limitations, their low computational requirements and ease of implementation make them suitable for quality improvement initiatives and resource-constrained settings [3], [6], [9]. Overall, rule-based and lexicon models can be understood as foundational yet context-limited approaches, highlighting their continued relevance for interpretability while underscoring the need for more adaptive and context-aware methods.

**Hybrid models**: The findings also highlight the continued relevance of hybrid and statistical machine-learning pipelines in sentiment analysis of EHR data. These approaches combine structured feature engineering with classifiers such as Support Vector Machines, Random Forests, and logistic regression, often incorporating techniques such as TF-IDF representations, embedding-enhanced models, and ensemble learning [6], [20], [29].

The synthesis indicates that such pipelines can improve semantic clustering and enhance the detection of minority sentiment classes compared with rule-based approaches, particularly on structured or semi-structured datasets [33]. This makes them valuable in contexts where computational efficiency and interpretability are prioritized. However, these methods remain dependent on extensive manual pre-processing and feature engineering, which limits scalability and adaptability. In contrast to transformer-based

models, their limited capacity for contextual understanding limits their ability to capture nuanced clinical narratives [3], [6]. As such, hybrid pipelines represent a transitional approach, offering performance gains over traditional methods while highlighting the trade-offs between efficiency, interpretability, and contextual depth.

The findings of this review are consistent with broader trends in clinical NLP research, where transformer-based models have become dominant due to their ability to capture contextual representations. However, unlike general NLP applications, sentiment analysis in EHRs introduces additional complexity related to clinical ambiguity, documentation bias, and ethical considerations. The limited representation of studies from the Global South further suggests that current advancements may not be readily transferable to low-resource healthcare settings, highlighting a critical gap in equitable digital health innovation.

**2) Frameworks used in the extraction and analysis of patient sentiment from Electronic Health Records:**

**Clinical informatics and decision support frameworks**: The review further highlights their role as key organizational enablers in integrating sentiment analysis into healthcare systems. These frameworks extend sentiment analysis beyond standalone analytics by embedding it within outcome measurement pipelines, oncology treatment extraction systems, and clinical decision support environments [27], [33]. This integration enables combining sentiment-derived signals with structured clinical data, supporting more comprehensive, context-aware predictive models. Importantly, the findings indicate that sentiment indicators can enhance risk stratification and the generation of real-world evidence, particularly when incorporated into predictive platforms that operate across multiple data sources [10], [34]. This suggests a shift toward more holistic clinical intelligence systems, where affective insights complement traditional clinical variables in informing care decisions. However, the effective implementation of these frameworks depends on organizational capabilities, particularly interoperability and platform infrastructure [16], [31]. The reliance on integrated systems underscores that technical feasibility alone is insufficient; institutional readiness and system-level coordination are critical to translating sentiment analytics into actionable clinical outcomes. Overall, these findings position sentiment analysis as an emerging component of digital clinical

governance, underscoring the need to strengthen alignment among predictive technologies, organizational infrastructure, and clinical decision-making processes.

**Patient-centric and experience mining frameworks**: The review also highlights the emergence of these frameworks, which position sentiment analysis as a mechanism to enhance patient engagement and service quality within healthcare systems [28], [32]. These frameworks incorporate approaches such as automated feedback analysis, user-centred EHR evaluation models, and predictive patient monitoring, enabling the extraction of affective insights that reflect patient experiences over time [8], [26]. Importantly, integrating sentiment data with demographic and behavioural information supports more personalized, context-aware interventions, aligning with broader trends in patient-centred data science [10], [11]. This integration allows healthcare providers to move beyond clinical indicators toward a more holistic understanding of patient needs and experiences. However, the synthesis suggests that, while these frameworks demonstrate promise, their implementation remains limited, particularly in translating insights into routine clinical practice. This highlights an ongoing challenge in operationalizing patient-centred analytics within existing healthcare infrastructures.

**Affective and social contextualization frameworks**: The findings highlight the growing importance of these frameworks in advancing sentiment analysis within healthcare systems. Unlike purely model-driven approaches, these frameworks situate sentiment extraction within broader socio-emotional and contextual analytics, enabling a more nuanced interpretation of patient data [19], [35]. In particular, the integration of emotional AI systems, IoT-enabled monitoring, and knowledge graph-based relational modelling reflects a shift toward multi-modal and context-aware architectures that extend beyond text-based analysis [12], [21]. This evolution is significant because it enables sentiment signals to be interpreted alongside social determinants of health, behavioural patterns, and disease trajectories, thereby enhancing predictive reasoning and clinical relevance [23], [33]. Rather than treating sentiment as an isolated feature, these frameworks position it within a relational, longitudinal understanding of patient states, which is critical for chronic disease management and personalized care. However, the synthesis also suggests that such integrative approaches remain unevenly developed and are often confined to experimental or domain-specific implementations. The practical integration of affective insights into real-time clinical workflows and decision support systems

remains limited, indicating a persistent gap between conceptual advances and operational deployment.

### 3) Algorithms/Techniques used in the extraction and analysis of patient sentiment from Electronic Health Records

**Lexicon- and rule-based implementation: The findings also reaffirm the foundational role of lexicon- and rule-based approaches in early healthcare sentiment analysis systems**. These methods, including lexicon-driven algorithms, ontology mapping, and rule-based pattern matching, have historically provided structured mechanisms for extracting sentiment from clinical text [13], [14], [28], [32]. Their reliance on predefined vocabularies and domain-specific rules makes them inherently transparent and interpretable, which is particularly important in clinical settings where explainability is a key requirement [6]. However, the synthesis indicates that these approaches are increasingly limited in their ability to capture contextual nuance and evolving clinical language. Their dependence on static lexicons and manually encoded rules limits adaptability, especially in complex narrative environments where meaning is highly context-dependent. As a result, while these techniques remain valuable for baseline systems and interpretable applications, they are often insufficient for more advanced, context-aware sentiment analysis tasks. This positions lexicon and rule-based methods as foundational but transitional approaches, highlighting their continued relevance for explainability while underscoring the need for more flexible and context-sensitive models in modern clinical NLP

**Transformer and attention models**: The findings highlight the dominance of transformer-based approaches in advancing sentiment analysis within EHR systems. By leveraging contextual embeddings, self-attention mechanisms, and multi-task learning, these models enable richer semantic representations of clinical narratives than earlier methods [24], [25],[36]. Techniques such as fine-tuning pre-trained clinical language models, analyzing attention weights, and handling long-form clinical text further enhance their ability to capture implicit affect and temporal patterns in patient records [6], [27], [30], [33]. This capability is particularly important in healthcare contexts, where sentiment is often subtle, context-dependent, and embedded within complex clinical documentation. In addition, integrating explainable AI techniques improves transparency

and interpretability, which are essential for clinical trust and regulatory acceptance [22]. Despite these strengths, the synthesis suggests that transformer-based models are still largely evaluated in controlled settings, with limited evidence of widespread clinical deployment. The findings further suggest that the dominance of transformer-based approaches is driven not only by performance improvements but also by their ability to capture contextual nuance in clinical narratives, a key limitation of traditional and rule-based methods. This highlights an ongoing gap between technical sophistication and real-world implementation in healthcare environments.

**Statistical ML and dimensionality reduction**: The review also identifies statistical machine learning approaches as an important intermediate stage in the evolution of sentiment analysis within healthcare. Techniques such as Support Vector Machines, Random Forests, Naïve Bayes, and logistic regression have been widely applied using TF-IDF and embedding-based features to model sentiment in clinical text [20], [29]. These approaches are often complemented by dimensionality reduction and feature selection techniques, including principal component analysis, to address the high dimensionality and sparsity of EHR data [9], [37]. The findings suggest that such methods offer practical advantages in computational efficiency and can improve classification performance, particularly for imbalanced datasets where minority sentiment classes are underrepresented [33]. This makes them suitable for structured environments with constrained computational resources. However, the synthesis indicates that these models are inherently limited by their reliance on manual feature engineering and their reduced capacity to capture contextual and semantic complexity. In contrast to transformer-based approaches, statistical models lack deep contextual awareness, which constrains their effectiveness in interpreting nuanced clinical narratives [25]. As a result, statistical machine learning methods can be understood as efficient but context-limited approaches that bridge the gap between early rule-based systems and more advanced deep learning models, while highlighting the trade-offs among interpretability, efficiency, and contextual understanding.

**4) How have application domains and use cases from sentiments enhanced patient-centric results in a healthcare context?**

The results section identified three major themes: clinical diagnosis and care optimization, patient experience and service evaluation, and specialized mental and

palliative care. Across the eleven primary studies examined, sentiment analytics was not viewed as a technical effort but rather as a clinical practice approach to improve decision-making, quality, and individualized care. The following discussion will critically examine the importance of each theme and its sub-cases in improving patient care.

**Clinical diagnostics and care optimization: The findings revealed that these areas were a prominent application domain in which sentiment extraction contributed to patient-centric outcome**s. Sub-application areas included cancer treatment identification, chronic disease risk assessment, goals-of-care documentation analysis, and AI-assisted decision support [3], [27], [30]. Evidence shows that combining sentiment information from EHR text with clinical data enhances predictive performance and facilitates personalized treatment recommendations [10], [38]. For instance, oncology-oriented NLP tools have shown that contextual evaluation of patient narratives improves treatment stratification and outcome prediction [6], [15]. Moreover, digital twin and predictive monitoring applications demonstrate the role of sentiment-driven analytics in anticipatory disease management [21], [33]. Frameworks for predictive analytics that combined text-based sentiment analysis with other structured variables enhanced personalized decision-making and post-COVID care optimization [10], [38]. There is evidence that these application examples improve anticipatory care, minimize documentation blind spots, and facilitate the development of real-world evidence [34], [40], thereby improving patient-centred clinical governance. Collectively, these applications make it clear that sentiment extraction improves clinical accuracy, anticipation, and the development of real-world evidence.

**Patient experience and service evaluation:** The findings also identified this area as a key focus, where Sentiment Analysis directly informs quality improvement strategies. Sub-cases included the automatic evaluation of patient feedback, virtual consultation analysis, user-focused EHR evaluation, and service optimization strategies [5], [28], [32]. Empirical evidence shows that NLP-based feedback analysis significantly transforms unstructured feedback into structured insights for improvement [13], [26]. Evidence also supports the idea that affective cue analysis in documentation improves insights into patient engagement and participation [2], [4]. Platform ecosystem research also supports the role of interoperability in facilitating scalable experience analysis in institutions [16], [31]. The

above evidence supports that sentiment-driven service evaluation improves responsiveness, transparency, and patient satisfaction.

**Specialized mental and palliative care: Specialized mental health and palliative care were identified as important areas where sentiment analysis has proven** highly beneficial for patient-centric outcomes. Sub-categorical areas included mental health occupation extraction, emotional state monitoring, palliative care support systems, and emotionally adaptive AI integration [14], [19], [25]. Evidence from NLP research in mental healthcare shows that free-text note analysis for contextual analysis helps early detection of emotional states and psychosocial determinants [9], [25]. In palliative settings, sentiment analysis-assisted monitoring helps in compassionate care and in anticipating symptoms [19], [27]. Research on the integration of emotional AI and IoT technology further exemplifies the adaptive nature of care models in response to patient emotions [35].

**5) What opportunities exist for leveraging sentiment analysis to enhance patient-centered care within EHR ecosystems?**

This section focuses on opportunities to strategically leverage sentiment analysis to improve patient-centric care in EHR systems. From the results section, three major opportunity themes emerged: decision support and risk management, care optimization and advanced phenotyping, and service refinement and psychosocial mapping. In the eleven studies included in the primary research, the opportunities were not presented as hypothetical developments but rather as scalable enhancements to existing NLP system deployments within interoperable, patient-centric information architectures [3], [33]. The following discussion critically assesses the opportunities identified in the results section, focusing on their respective sub-opportunities.

**Clinical decision support and risk management**: The findings have highlighted the importance of improved predictive decision-support tools, the early detection of risks from unstructured clinical narratives, and the incorporation of sentiment markers into CDS dashboard analytics. Evidence from cancer and chronic disease models has shown that including affective markers can improve predictive performance beyond structured data alone [3], [21]. The automatic detection of goal-of-care discussions also highlights the use of sentiment-informed markers as quantifiable quality and safety metrics [27].

Secondary analyses have reinforced the importance of integrating NLP outputs into CDS systems to improve anticipatory care and develop real-world evidence [10], [31]. In this way, sentiment analysis can be seen as an opportunity to improve proactive risk management and personalized intervention planning in EHR-supported settings.

**Care optimization and advanced phenotyping**: The findings also emphasized advanced disease phenotyping, cohort stratification based on narrative embeddings, and digital twin integration as important sub-opportunities [20], [37]. Embedding-based sentiment analytics improves the identification of latent psychosocial and behavioural phenotypes encoded in unstructured EHR narratives [6], [25]. Evidence from automated disease cohort identification indicates that contextual word embeddings enhance the classification of intricate diseases [20], [29]. The integration of these findings into predictive monitoring systems and digital twin environments facilitates the dynamic modelling of patient pathways [11], [21]. Secondary literature also supports the finding that these strategies correspond with the principles of precision medicine and care optimization [34], [39]. Therefore, sentiment-driven phenotyping is a methodological extension of patient profiling from structured diagnostics.

**Service optimization and psychosocial mapping: The findings also showed potential for service optimiz**ation through automated feedback mining, document quality evaluation, and psychosocial determinant mapping [2], [32]. Sub-opportunities were identified as extracting occupation- and social-contextual features from mental health documents and evaluating patient experience comments for institutional learning [14], [28]. Evidence indicates that sentiment-driven service analytics improves responsiveness to communication deficits and structural inequities [5], [13]. In addition, the combination of emotional AI and smart monitoring technology ensures adaptive healthcare systems that are responsive to patients' affective experiences [19], [35]. These findings demonstrate that sentiment analysis extends beyond diagnosis to broader psychosocial mapping and transformation.

**6) What challenges exist for the reliability, generalizability, or applicability of sentiment models for EHRs?**

The challenge domains identified in the results section are linguistic and semantic complexity, data governance and technical constraints, and operational noise and

workflow variability. In the eleven primary studies identified in this research, the challenges are observed rather than inferred, indicating systemic shortcomings in model robustness. The discussion section below will critically discuss the challenge domains identified in the results section.

**Linguistic and semantic complexity**: The results have identified contextual ambiguity, negations, variations in clinical jargon, and implicit affective communication as sub-problems considered significant for ensuring model reliability. The clinical text is usually associated with hedging, abbreviations, and temporal narratives, which can result in reduced accuracy in sentiment classification [2], [3]. NLP studies in oncology and psychiatry have shown that certain words can have opposite sentiment polarities, making it difficult to generalise models across domains [15], [30]. Systematic reviews have found that even transformer models, which are considered more accurate, can be affected by domain shift and semantic drift, even when tested on external datasets [40],[41].

**Data governance and technical constraints**: The findings also emphasized the importance of privacy-preserving issues, de-identification challenges, the lack of annotated datasets, and interoperability issues as sub-challenges. Free-text EHR data are highly vulnerable to re-identification, thereby hindering data sharing for model validation and generalization testing [3], [6]. Moreover, studies have shown a lack of annotated datasets for sentiment analysis in the healthcare domain, thereby increasing the need for small-scale institutional datasets with low external validity [9], [25]. A platform ecosystem analysis identifies data silos and a lack of documentation standards as technical hurdles to large-scale implementation of NLP solutions [16], [34].

**Operational noise**: The findings also highlighted that documentation inconsistencies, clinician workload pressures, and EHR interface variability are sub-challenges contributing to operational noise. Variability in how clinicians record patients' emotions or experiences introduces systematic bias into sentiment detection outputs [5], [32]. Research studies on user-centric EHR evaluation show that documentation is influenced by usability constraints rather than by the need for expressive completeness, thereby restricting the richness of sentiment signals [26]. Moreover, real-world studies on decision aid evaluations have shown that AI models tend to perform poorly when

compared to controlled research environments, owing to workflow misalignment [33], [38].

Overall, from a governance perspective, the findings indicate that sentiment analysis in EHRs is evolving faster at the technological level than at the regulatory and organizational levels. This imbalance raises concerns regarding interpretability, accountability, and clinical integration, suggesting the need for frameworks that explicitly incorporate ethical, legal, and socio-technical considerations into model deployment.

### 3.11. Synthesis of Key Limitations and Framework Motivation

The synthesis of the eleven included studies reveals several recurring limitations that motivate the development of a conceptual framework. First, existing research is predominantly technically oriented, with a strong emphasis on transformer-based models and hybrid pipelines, yet limited integration within interoperable EHR ecosystems. Second, the generalisability of the proposed approaches remains constrained by reliance on institution-specific datasets, domain-shift challenges, and limited cross-context validation. Third, although sentiment analysis techniques show promise for extracting psychosocial and affective insights from clinical narratives, their application to real-time clinical decision support and workflow integration remains underdeveloped. Finally, governance-related challenges, including data de-identification, annotation inconsistencies, algorithmic bias, and fragmented data environments, continue to hinder scalability, regulatory alignment, and practical implementation. Collectively, these limitations highlight the need for an integrated socio-technical perspective that connects methodological advancements with organizational, clinical, and governance considerations. Building on the identified methodological patterns and implementation gaps, the study advances a conceptual framework that integrates these insights into a structured socio-technical perspective.

### 3.12. Socio-Technical Sentiment Governance Framework (STSGF)

The Socio-Technical Sentiment Governance Framework (STSGF) was derived through an inductive thematic synthesis of the included studies. Specifically, recurring themes related to technological approaches (e.g., transformer models), organizational factors (e.g., clinical workflows), and governance considerations (e.g., ethics, privacy, and interpretability) were identified and grouped into higher-level conceptual categories.

These categories were then integrated into a unified framework to reflect the socio-technical nature of sentiment analysis in EHR systems. Building on the cross-study synthesis, this study proposes the Socio-Technical Sentiment Governance Framework (STSGF) (Figure 4) as an integrative model for understanding and guiding the application of sentiment analysis within Electronic Health Record (EHR) systems. The framework is developed in response to the identified limitations in the literature, particularly the fragmentation between technological innovation, clinical application, and governance considerations.

The proposed conceptual framework captures the multi-dimensional nature of sentiment analysis within Electronic Health Record (EHR) systems by illustrating how technical, clinical, and governance-related elements interact to support patient-centred care. At its core, the framework reflects the study's synthesis of the literature, which shows that sentiment analysis is not a purely computational task, but a socio-technical capability embedded within complex healthcare environments. The framework emphasizes the role of advanced natural language processing techniques, particularly transformer-based models, in extracting affective and contextual insights from unstructured clinical narratives. However, it also highlights that underlying data characteristics, including data quality, annotation consistency, and contextual variability across institutions, shape the effectiveness of these approaches. These data-related considerations influence both model performance and the reliability of derived insights.

In addition, the framework illustrates how sentiment outputs are intended to inform clinical and operational processes, such as patient monitoring, care coordination, and decision support. Despite this potential, the synthesis indicates that real-world integration into clinical workflows remains limited, suggesting a gap between methodological capability and practical application. The framework further incorporates governance considerations, including privacy, ethical use, interpretability, and regulatory compliance. These factors are critical in ensuring that sentiment analysis systems are trustworthy, accountable, and aligned with healthcare standards. Overall, the framework provides an integrated perspective that connects methodological innovation with clinical utility and governance requirements, thereby offering a structured lens for understanding the current state and future direction of sentiment analysis in EHR systems. The proposed framework should therefore be interpreted as a synthesis-driven

conceptual model rather than a validated implementation architecture, providing a structured lens to guide future empirical development.

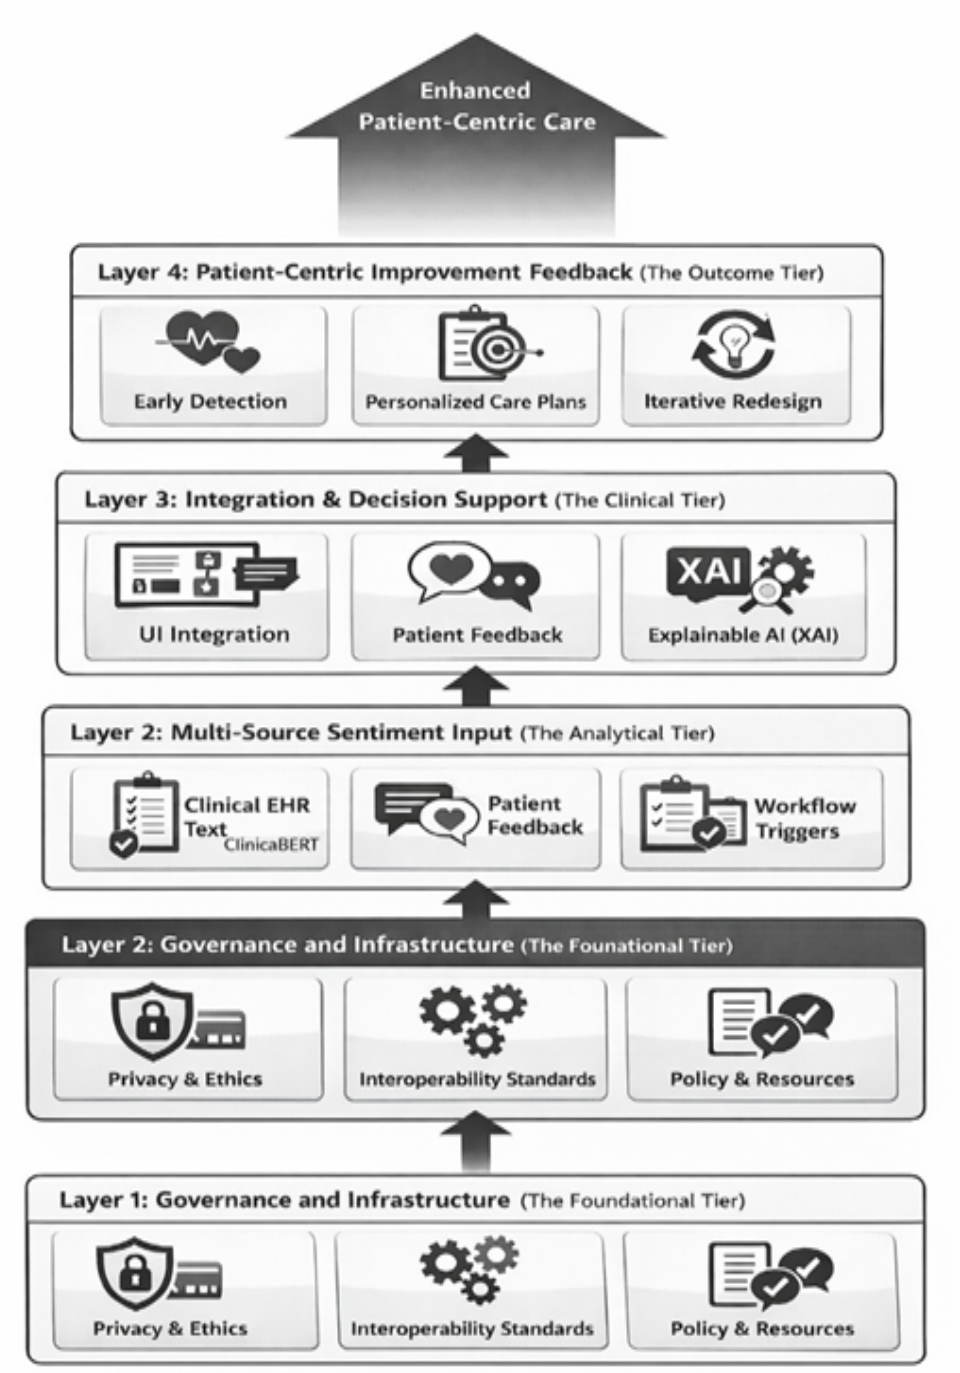



**Figure 4.** Socio-Technical Sentiment Governance Framework (STSGF)

### 3.13. Implications

The findings of this study have important implications for research, practice, and policy in the evolving field of sentiment analysis within Electronic Health Record (EHR) systems.

From a research perspective, the review highlights a critical shift from model-centric development toward integrated socio-technical approaches, aligning with broader trends in clinical natural language processing and digital health analytics. While transformer-based models demonstrate strong performance, future research should prioritize cross-institutional validation, standardized annotation protocols, and context-aware modelling, particularly to address challenges related to domain shift and dataset bias. In addition, the growing emphasis on explainable artificial intelligence (XAI) reflects an emerging requirement for transparency and trust in clinical decision-making environments.

From a practical standpoint, the findings reinforce that the value of sentiment analysis lies not only in predictive performance but in its operational integration within clinical workflows. Healthcare organizations should therefore prioritize incorporating sentiment-derived insights into decision support systems, patient monitoring, and care coordination processes. This requires alignment among technological capabilities, clinician engagement, and organizational readiness to support a transition from experimental applications to clinically actionable intelligence.

At the policy and governance level, the study underscores the need for robust regulatory and ethical frameworks to guide the deployment of sentiment analysis in healthcare. Persistent concerns about data privacy, algorithmic bias, and accountability underscore the importance of governance mechanisms that ensure compliance with existing regulatory standards while accommodating emerging AI-driven tools. Policymakers should therefore focus on developing adaptive guidelines that support responsible, transparent, and equitable implementation across diverse healthcare contexts. Collectively, this study advances the understanding of sentiment analysis in EHRs by positioning it as a data-driven, socio-technical capability rather than a standalone analytical task. The implications point to a future research agenda centred on scalability, interoperability, and governance alignment, all of which are essential to realizing the full potential of sentiment analysis in patient-centred healthcare.

## 4. CONCLUSION

This article summarises the history of sentiment analysis in Electronic Health Records (EHRs), with special reference to the impact of the introduction of transformers and

clinical NLP. The effectiveness of the technology in this field, however, as the article reveals, remains limited to the domain or environment in which it was developed. The effectiveness of sentiment analysis in this field also largely depends on various critical socio-technical factors. The critical barriers to the effectiveness of sentiment analysis in this field include dataset, annotation, and domain biases, as well as a lack of geographic diversity. In this respect, by situating sentiment analysis as a socio-technical capability, this research addresses the gap between technological innovation and healthcare processes, thereby aligning with emerging trends in digital health that highlight the importance of NLP integration into clinical governance structures. In terms of future directions, this research recognises the need for validation, benchmarking, and the development of interpretable models as vital steps for building trust among clinical professionals. In addition, future work will need to focus on integrating sentiment-aware information into decision support systems. In this respect, advancing sentiment analysis as a technology to improve healthcare processes will require the development of flexible regulatory and ethical frameworks that facilitate the effective transfer of analytical innovations into patient-centric improvements in the overall patient experience.

## ACKNOWLEDGMENT

The authors wish to express their gratitude to the academic and research communities that have contributed to the knowledge base that enabled them to conduct a systematic literature review. We are thankful to the reviewers/editors who have provided valuable insights that have helped to enhance the quality of this manuscript.

## REFERENCES

[1] M. Madar Bandu, S. Thirkateh, K. Sai, S. Reddy, P. M. Reddy, and V. Rishikesh, "Electronic Health Records Platform: Patient-Centric Healthcare Management," *Alinteri J. Agr. Sci*, vol. 39, no. 2, p. 39, 2021.

[2] J. M. Butler *et al.*, "Clinician documentation of patient centered care in the electronic health record," *BMC Med. Inform. Decis. Mak.*, vol. 22, no. 1, pp. 1–12, 2022, doi: 10.1186/s12911-022-01794-w.

[3] L. Wang *et al.*, "Assessment of Electronic Health Record for Cancer Research and Patient Care Through a Scoping Review of Cancer Natural Language Processing," *JCO Clin. Cancer Informatics*, no. 6, pp. 1–17, 2022, doi: 10.1200/cci.22.00006.

[4] K. Liosatos, G. Tobiano, and B. M. Gillespie, "Patient participation in surgical wound care in acute care settings: An integrative review," *Int. J. Nurs. Stud.*, vol. 157, p. 104839, 2024, doi: 10.1016/j.ijnurstu.2024.104839.

[5] Y. V. Shi and S. Komiak, "Unveiling patient-centric interactions in virtual consultation: A comprehensive text mining approach," *Health Informatics J.*, vol. 31, no. 1, 2025, doi: 10.1177/14604582251327093.

[6] P. Thatoi, A. Shiwlani, H. Qureshi, R. Choudhary, H. A. Qureshi, and S. Kumar, "Natural Language Processing (NLP) in the Extraction of Clinical Information from Electronic Health Records (EHRs) for Cancer Prognosis," *Int. J. Membr. Sci. Technol.*, vol. 10, no. 4, pp. 2676–2694, 2023.

[7] S. Jadhav, B. Shiragapur, R. Purohit, and N. Jha, "Sentiment analysis in public health: a systematic review of the current state, challenges, and future directions," *Front. Public Heal.*, vol. 12, no. 2024, pp. 0–3, 2024, doi: 10.3389/fpubh.2024.1435211.

[8] S. M. Sivalingam and S. Thisin, "A new framework to enhance healthcare monitoring using patient-centric predictive analysis," *Int. J. Electr. Comput. Eng.*, vol. 14, no. 3, pp. 3295–3302, 2024, doi: 10.11591/ijece.v14i3.pp3295-3302.

[9] Md Russel Hossain, Shohoni Mahabub, Abdullah Al Masum, and Israt Jahan, "Natural Language Processing (NLP) in Analyzing Electronic Health Records for Better Decision Making," *J. Comput. Sci. Technol. Stud.*, vol. 6, no. 5, pp. 216–228, 2024, doi: 10.32996/jcsts.2024.6.5.18.

[10] A. Zahid and R. Sharma, "Personalized Health Care in a Data-Driven Era: A Post–COVID-19 Retrospective," *Mayo Clin. Proc. Digit. Heal.*, vol. 1, no. 2, pp. 162–171, 2023, doi: 10.1016/j.mcpdig.2023.04.002.

[11] M. Amoei and D. Poenaru, "Patient-centered data science: an integrative framework for evaluating and predicting clinical outcomes in the digital health era," *arXiv Prepr. arXiv2408.02677*, no. Ml, pp. 1–12, 2024.

[12] H. S. Al Khatib *et al.*, "Patient-centric knowledge graphs: a survey of current methods, challenges, and applications," *Front. Artif. Intell.*, vol. 7, no. October, pp. 1–21, 2024, doi: 10.3389/frai.2024.1388479.

[13] E. Vashishtha and H. Kapoor, "Enhancing Patient Experience by Automating and Transforming Free Text into Actionable Consumer Insights: A Natural Language Processing (NLP) Approach," *Int. J. Heal. Sci. Res.*, vol. 13, no. 10, pp. 275–288, 2023, doi: 10.52403/ijhsr.20231038.

[14] N. Chilman *et al.*, "Text mining occupations from the mental health electronic health record: A natural language processing approach using records from the Clinical Record Interactive Search (CRIS) platform in south London, UK," *BMJ Open*, vol. 11, no. 3, pp. 1–11, 2021, doi: 10.1136/bmjopen-2020-042274.

[15] A. Bhatia *et al.*, "Application of Natural Language Processing in Electronic Health Record Data Extraction for Navigating Prostate Cancer Care: A Narrative Review," *J. Endourol.*, vol. 38, no. 8, pp. 852–864, 2024, doi: 10.1089/end.2023.0690.

[16] M. C. Chibuike, S. S. Grobbelaar, and A. Botha, "Overcoming Challenges for Improved Patient-Centric Care: A Scoping Review of Platform Ecosystems in Healthcare," *IEEE Access*, vol. 12, no. January, pp. 14298–14313, 2024, doi: 10.1109/ACCESS.2024.3356860.

[17] S. Arunprasath and S. Annamalai, *Improving patient centric data retrieval and cyber security in healthcare: privacy preserving solutions for a secure future*, vol. 83, no. 27, doi: 10.1007/s11042-024-18253-5. Springer US.

[18] H. Chamboko and B. Ndlovu, "Twitter ( X ) Sentiment Analysis on Monkeypox : A Systematic Literature Review," *Int. J. Informatics Dev.*, vol. 14, no. 2, pp. 629–639, 2025, doi: 10.14421/ijid.2025.5196.

[19] K. Maguraushe and B. M. Ndlovu, "The use of smart technologies for enhancing palliative care: A systematic review," *Digit. Heal.*, vol. 10, 2024, doi: 10.1177/20552076241271835.

[20] B. S. Glicksberg, R. Miotto, K. W. Johnson, K. Shameer, and L. Li, "Automated disease cohort selection using word embeddings from Electronic Health Records," *Pacific Symp. Biocomput.*, pp. 145–156, 2018.

[21] S. Hadebe, B. Ndlovu, and K. Maguraushe, "Managing Diabetes Using Machine Learning and Digital Twins," *Indones. J. Innov. Appl. Sci.*, vol. 5, no. 2, pp. 145–162, 2025, doi: 10.47540/ijias.v5i2.1981.

[22] B. Ndlovu, K. Maguraushe, and O. Mabikwa, "Machine Learning and Explainable AI for Parkinson's Disease Prediction: A Systematic Review," *Indones. J. Comput. Sci.*, vol. 14, no. 2, pp. 2357–2386, 2025, doi: 10.33022/ijcs.v14i2.4837.

[23] A. L. Vladimir, "Natural Language Processing for Clinical Text Analysis and Patient

Outcome Prediction : Unlocking Unstructured Data in Electronic Health Records," *ResearchGate*, 2026.

[24] T. Classification, "SS symmetry Balanced Knowledge Transfer in MTTL-ClinicalBERT : A Symmetrical Multi-Task Learning Framework for Clinical," *IEEE Access*, vol. 17, no. 6, pp. 1–26.

[25] G. Delanerolle, Y. Bouchareb, S. Shetty, H. Cavalini, and P. Phiri, "A Pilot Study Using Natural Language Processing to Explore Textual Electronic Mental Healthcare Data," *Informatics*, vol. 12, no. 1, 2025, doi: 10.3390/informatics12010028.

[26] R. Sreejith and K. R. Sinimole, "User-centric evaluation of EHR software through NLP-driven investigation: Implications for product development and user experience," *J. Open Innov. Technol. Mark. Complex.*, vol. 10, no. 1, p. 100206, 2024, doi: 10.1016/j.joitmc.2023.100206.

[27] R. Y. Lee *et al.*, "Assessment of Natural Language Processing of Electronic Health Records to Measure Goals-of-Care Discussions as a Clinical Trial Outcome," *JAMA Netw. Open*, vol. 6, no. 3, p. E231204, 2023, doi: 10.1001/jamanetworkopen.2023.1204.

[28] P. Mathur *et al.*, "Automated analysis of ambulatory surgery patient experience comments using artificial intelligence for quality improvement: A patient centered approach," *Intell. Med.*, vol. 5, p. 100043, 2021, doi: 10.1016/j.ibmed.2021.100043.

[29] T. E. Brunekreef *et al.*, "Text Mining of Electronic Health Records Can Accurately Identify and Characterize Patients With Systemic Lupus Erythematosus," *ACR Open Rheumatol.*, vol. 3, no. 2, pp. 65–71, 2021, doi: 10.1002/acr2.11211.

[30] J. Zeng *et al.*, "Natural Language Processing to Identify Cancer Treatments With Electronic Medical Records," *JCO Clin. Cancer Informatics*, no. 5, pp. 379–393, 2021, doi: 10.1200/cci.20.00173.

[31] M. A. Saberi, H. Mcheick, and M. Adda, "From Data Silos to Health Records Without Borders: A Systematic Survey on Patient-Centered Data Interoperability," *Inf.*, vol. 16, no. 2, pp. 1–31, 2025, doi: 10.3390/info16020106.

[32] M. Khanbhai, P. Anyadi, J. Symons, K. Flott, A. Darzi, and E. Mayer, "Applying natural language processing and machine learning techniques to patient experience feedback : a systematic review," 2021, doi: 10.1136/bmjhci-2020-100262.

[33] R. Stewart, J. Chaturvedi, and A. Roberts, "Natural language processing – relevance to patient outcomes and real-world evidence," *Expert Rev. Pharmacoecon. Outcomes Res.*, vol. 24, no. 1, pp. 5–10, 2024, doi: 10.1080/14737167.2023.2275670.

[34] F. S. Almalki *et al.*, "Leveraging Big Data Analytics To Improve Healthcare Quality : A,"

vol. 10, no. 1, 2025.

[35] N. Bout, G. Moukhliss, H. Belhadaoui, N. Afifi, and M. Abik, "Integrating emotional AI, IoT, and robotics for patient-centered healthcare: challenges and future directions," *Discov. Internet Things*, vol. 5, no. 1, 2025, doi: 10.1007/s43926-025-00174-w.

[36] J. Machemedze and B. Ndlovu, "Transformer-Based Models for Electronic Health Records and Omics in Healthcare : A Systematic Literature Review," *J. Appl. Informatics Comput.*, vol. 10, no. 1, pp. 90–105, 2026,doi: 10.30871/jaic.v10i1.11893.

[37] N. Hiob and S. Lessmann, "Health Analytics : a systematic review of approaches to detect phenotype cohorts using electronic health records," *J. Am. Med. Informatics Assoc.*, vol. 23 no.6, no. pp. 1157–1171, 2016, doi: 10.1093/jamia/ocw018, 2016.

[38] J. Cochrane, O. L. Roberts, K. Ribbons, B. Y. Tan, and M. Nilsson, "Exisiting Ai Driven Patient-Based Predictive Decision Aids Which Inform Knee Osteoarthritis Interventions: a Scoping Review," *Osteoarthr. Cartil.*, vol. 33, no. 2025, pp. S102–S103, 2025, doi: 10.1016/j.joca.2025.02.139.

[39] G. Lyu, "Data-driven decision making in patient management : a systematic review," *BMC Med. Inform. Decis. Mak.*, vol. 25(1), no. 239. doi: 10.1186/s12911-025-03072–x, 2025.

[40] O. T. Chikumo and B. Ndlovu, "Transformer-based Models for Cardiovascular Disease Predictions from Electronic Health Records : A Systematic Review," *J. Appl. Informatics Comput.*, vol. 10, no. 1, 2026, doi:10.30871/jaic.v10i1.11899

[41] S. S. Sibanda and B. Ndlovu, "Explainable Transformer and Machine Learning Models in Predicting Tuberculosis Treatment Outcomes . A Systematic Review," *J. Appl. Informatics Comput.*, vol. 10, no. 1, pp. 150–164, 2026 ,doi: 10.30871/jaic.v10i1.11846