

Integrating Natural Language Processing into Electronic Health Records for Automated Clinical Decision Support

Ke Liu¹, Shuang Hu², Dandan Xu^{3,*}

¹ Xuzhou City traditional Chinese medicine hospital, Jiangsu221000, China

² Xiangya School of Nursing, Central South University

³ Department of Emergency intensive care unit, The Affiliated Hospital of Xuzhou Medical University, Xuzhou, China

*Email: 214731522@qq.com

Abstract

This study explores the integration of Natural Language Processing (NLP) into Electronic Health Records (EHRs) to enhance Clinical Decision Support Systems (CDSS). Despite the vast amount of data in EHRs, much of the unstructured data, such as clinical notes, remains underutilized. NLP can extract meaningful information from this unstructured text, bridge the gap between data availability and actionable clinical knowledge. We developed an NLP pipeline using the MIMIC-III dataset, to achieve high performance in entity recognition and relation extraction. The integrated CDSS demonstrated improved accuracy in key clinical tasks, such as early sepsis detection and in-hospital mortality prediction. The system maintained real-time responsiveness, which generated patient-specific recommendations within 3.5 seconds. This study highlights the potential of NLP to transform CDSS by leveraging unstructured clinical data, ultimately improving patient outcomes and healthcare efficiency.

Keywords: Natural Language Processing, Electronic Health Records, Clinical Decision Support Systems, Healthcare Informatics

Article History:

Received July 05, 2024

Revised December 20, 2024

Accepted November 12, 2024

Available Online December 10, 2024

Integrating Natural Language Processing into Electronic Health Records for Automated Clinical Decision Support

Abstract

This study explores the integration of Natural Language Processing (NLP) into Electronic Health Records (EHRs) to enhance Clinical Decision Support Systems (CDSS). Despite the vast amount of data in EHRs, much of the unstructured data, such as clinical notes, remains underutilized. NLP can extract meaningful information from this unstructured text, bridge the gap between data availability and actionable clinical knowledge. We developed an NLP pipeline using the MIMIC-III dataset, to achieve high performance in entity recognition and relation extraction. The integrated CDSS demonstrated improved accuracy in key clinical tasks, such as early sepsis detection and in-hospital mortality prediction. The system maintained real-time responsiveness, which generated patient-specific recommendations within 3.5 seconds. This study highlights the potential of NLP to transform CDSS by leveraging unstructured clinical data, ultimately improving patient outcomes and healthcare efficiency.

Keywords: Natural Language Processing, Electronic Health Records, Clinical Decision Support Systems, Healthcare Informatics

1. Introduction

Electronic Health Records (EHRs) have become a fundamental component of modern healthcare, which provides a comprehensive digital repository of patient information (Hossain et al., 2024). With the increasing digitization of healthcare systems globally, EHRs now contain vast amounts of rich, longitudinal data that, in theory, could significantly enhance clinical decision-making and patient outcomes (Warbhe & Verma, 2024). However, much of the information stored within EHRs remains underutilized (Sivarethinamohan et al., 2021). The unstructured portions of EHRs harbor valuable insights that traditional data retrieval methods often fail to capture efficiently (Sim et al., 2024; Sim et al., 2023). Structured fields like diagnosis codes and lab results are more easily queried and analyzed, but they represent only a fraction of the knowledge embedded within the complete record (Putra et al., 2019; Mellia et al., 2020). The underexploitation of unstructured data points to a critical gap between data availability and actionable clinical knowledge (Afshar et al., 2023). Addressing this gap is essential for realizing the full potential of EHR systems (Yadav et al., 2023). Natural Language Processing (NLP) has emerged as a powerful technology for extracting meaningful information from unstructured text and offering a transformative solution to the challenges of EHR data utilization (Chaudhry et al., 2018). NLP enables the interpretation of clinical narratives, identification of relevant medical concepts, extraction of temporal information, and synthesis of patient trajectories over time (Despotou et al., 2018). Recent advances in deep learning and transformer-based architectures have further expanded the capabilities of NLP systems, making it possible to capture nuanced clinical contexts that were previously inaccessible. In

ISSN © 2024 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author.

See: <https://inatgi.in/index.php/jaihbe/index> for more information. DOI: 10.63646/AKKZ8260

the realm of clinical decision-making, NLP can serve multiple roles (Yuvaraj et al., 2021): identifying patients eligible for specific treatments or clinical trials, flag potential diagnostic errors, predicting patient deterioration, and support personalized medicine initiatives (Cruz et al., 2019). By bridging the divide between unstructured data and structured analysis, NLP empowers clinicians with richer, more timely insights, ultimately enhancing the quality, safety, and efficiency of healthcare delivery (Senders, et al., 2020).

Clinical Decision Support Systems (CDSS) have long been recognized as essential tools in improving healthcare outcomes by providing evidence-based recommendations at the point of care (Houssein et al., 2021). Traditional CDSS models often rely heavily on rule-based logic and structured data inputs (Afshar et al., 2022; Sedgwick et al., 2023). The integration of NLP technologies into CDSS design represents a significant evolution toward "smart" CDSS—systems that are adaptive, context-aware, and capable of learning from vast and heterogeneous clinical datasets. Smart CDSS can analyze real-time patient narratives, extract actionable insights from prior clinical encounters, and provide recommendations that account for both structured and unstructured information (Jerfy et al., 2024). This new generation of decision support tools holds the promise of reducing cognitive burden on healthcare providers, and promoting personalized care pathways (Hossain et al., 2023). The successful deployment of NLP-enabled CDSS could serve as a critical enabler for broader initiatives in precision medicine and value-based care models.

2.Literature Review

Over the past decade, significant advances in medical text mining have transformed how unstructured clinical data can be harnessed for meaningful insights (Clay, et al., 2025). Central to these advances are key natural language processing tasks such as Named Entity Recognition (NER), relation extraction, and text classification (Tu et al., 2023a). NER aims to automatically identify and categorize key medical concepts—such as diseases, medications, procedures, and anatomical locations—from clinical narratives (Sousa et al., 2023). Early rule-based NER systems have gradually been supplanted by machine learning and, more recently, deep learning approaches, with models like BiLSTM-CRF and transformer-based architectures such as BERT and BioBERT demonstrating superior performance across diverse clinical datasets (Tu et al., 2023b). Beyond merely recognizing entities, relation extraction techniques seek to identify and structure the relationships among clinical concepts (Katyan, Gulati, & Upreti, 2024). Effective relation extraction is critical for constructing knowledge graphs and supporting inferential reasoning in clinical applications (Sousa et al., 2023). Text classification methods are widely used to categorize clinical documents into diagnostic groups, predict disease progression, or assess the severity of patient conditions (Jamaluddin & Wibawa, 2021). With the emergence of large annotated corpora like MIMIC-III and publicly available benchmarks, researchers have been able to fine-tune increasingly sophisticated models for specific clinical tasks, push the boundaries of what can be achieved in medical text mining (Pai, et al., 2014; Senthil, et al., 2024). These advances not only enable a deeper understanding of clinical narratives but also lay the groundwork for more intelligent and context-sensitive decision support systems (Burgt et al., 2022). Parallel to the evolution of medical text mining, Clinical Decision Support Systems (CDSS) have undergone several distinct development stages (Shahsavarani et al., 2015). The earliest generation of CDSS emerged in the 1970s and

1980s, typified by expert systems like MYCIN and INTERNIST-I, which used manually encoded rules and heuristics to provide recommendations (Musen et al., 2021). These rule-based systems, while groundbreaking at the time, faced challenges related to scalability, maintenance, and the brittleness of their inference mechanisms. As healthcare data became more digitized in the 1990s and 2000s, the second generation of CDSS focused on integrating structured EHR data with evidence-based guidelines, often relying on if-then rules and clinical pathways to drive decision support (Afshar, et al., 2019). While these systems improved clinical workflow integration, they still struggled to adapt to the variability and complexity of real-world clinical environments (Kumbhakarna, Kulkarni, & Dhawale, 2020). In the most recent phase, driven by advances in machine learning and data-driven approaches, CDSS has entered a new era of intelligent systems capable of learning from historical data, adapting to new clinical evidence, and incorporating both structured and unstructured data inputs (Afshar, et al., 2022). Contemporary smart CDSS leverage predictive modeling, natural language understanding, and dynamic updating mechanisms to provide more personalized, context-aware, and continuously evolving recommendations (Karthik et al., 2025). The integration of NLP into CDSS design represents a natural convergence of these trends, offer the possibility of systems that can understand free-text clinical notes, reason across patient histories, and generate nuanced, patient-specific guidance in real time (Mahire, Saravanan, Murali, & P. N., 2024; Jenders, 2024). This trajectory highlights the increasing sophistication of both underlying technologies and the ambition of CDSS applications to move beyond static rule-based alerts toward truly intelligent clinical assistants (Jenders, 2017).

3. Methods

3.1 Data Sources (MIMIC-III)

The primary dataset for this study was the Medical Information Mart for Intensive Care III (MIMIC-III), a publicly available critical care database containing de-identified EHR data. MIMIC-III comprises 53,423 distinct adult hospital admissions to five Intensive Care Units (ICUs) at the Beth Israel Deaconess Medical Center between 2001 and 2012 (Ridoy et al., 2024). This rich dataset includes patient demographics, structured data, and unstructured clinical notes (Qu & Li, 2024). For our purposes, we extracted the relevant clinical narrative text along with corresponding patient identifiers and timeline data.

We selected all notes associated with adult patient stays to ensure a uniform cohort. Patients with ICU stays shorter than 24 hours or missing key demographic information were excluded to avoid sparsely documented cases. The final sample consisted of several thousand ICU stays, which encompassed the full variety of note types. All processing was performed in accordance with the data use agreement and ethical guidelines for de-identified data.

Initial preprocessing of the raw data involved several steps to prepare it for NLP analysis. We consolidated notes at the patient-episode level, and merged notes chronologically for each ICU stay to create a comprehensive narrative timeline per admission. Duplicate notes were removed, and text encoding issues were resolved. We partitioned the dataset at the patient level into training, validation, and test subsets, to prevent information leakage across splits. 70% of patients were assigned to the training set, 15% to the validation set, and 15% to a held-out test set. This split ensured that no patient's data appeared in more than one partition. Stratification

by ICU unit and patient age was employed to maintain a balanced distribution of clinical scenarios across splits.

Subsequent preprocessing of the narrative text incorporated common cleaning routines to normalize and structure the free-text data. This included lowercasing all text and deleting residual de-identification placeholders that are present in MIMIC-III. Punctuation was mostly retained except for rare symbols; commas and periods were preserved to maintain sentence boundaries, but unusual unicode characters and non-alphanumeric symbols were discarded. Numeric tokens were normalized by replacing explicit numbers with a placeholder token to reduce lexical sparsity. We also expanded frequent medical abbreviations and acronyms using a custom dictionary derived from clinical ontologies to improve consistency in entity recognition (Cui et al., 2023). The cleaned text was segmented into sentences using a rule-based sentence splitter designed for clinical text and tokenized into word units using the spaCy toolkit. These preprocessing steps produced cleaned, tokenized text ready for downstream NLP analysis.

3.2 NLP Pipeline

The core of our methodology is an end-to-end NLP pipeline that processes cleaned clinical text to extract structured information, which is made available to the Clinical Decision Support System (CDSS). The pipeline consists of three main components: Text Cleaning, Entity Recognition, and Relation Modeling. Each component is described below, along with details of its design and implementation.

3.2.1 Text Cleaning

After the initial data preprocessing above, we applied a focused text-cleaning procedure on each clinical note to optimize it for information extraction (Guyen & Lamurias, 2023). Specifically, we performed the following steps:

(a) Token normalization: All tokens were lowercased, and numeric values were replaced with the token <num>. We also removed obvious patient identifiers or placeholders to avoid spurious correlations.

(b) Punctuation and symbol handling: We preserved sentence-terminating punctuation to maintain sentence boundaries but removed most other punctuation except where they were likely part of clinical abbreviations. We replaced any non-ASCII or unusual Unicode characters with standard equivalents.

(c) Abbreviation expansion: We applied a clinical abbreviation dictionary to expand common shorthand based on the UMLS Metathesaurus and curated lists. Abbreviations with multiple possible expansions were disambiguated using the surrounding context in the sentence.

(d) Stopword and short-token filtering: Common English stop words were not removed globally, since many could be part of medical terms; however, we did filter out isolated single-character tokens which tended to represent artifacts.

(e) Lemmatization: We used *specie*'s lemmatized to reduce tokens to their lemma forms, improving consistency. However, care was taken to not remove clinical meaning.

(f) Sentence segmentation: We employed a clinical rule-based splitter to divide notes into sentences. This ensured that downstream entity recognition and relation extraction could work at the sentence level.

These cleaning procedures standardized the text, reduced noise, and produced consistent tokenized sentences, which served as input for the NER and relation extraction models (Zhou et al., 2023).

3.2.2 Entity Recognition

The next step was to identify clinically relevant entities within the text. We defined entity categories of interest that would be most useful for decision support. These categories included diagnoses/conditions, signs or symptoms, medications, laboratory tests, procedures, and anatomical locations. A small set of additional utility categories was also tagged. To train and evaluate the entity recognizer, we created a labeled dataset as follows:

Annotation strategy: A corpus of 2,000 representative sentences was randomly sampled from the training set notes and manually annotated by two clinical experts. Annotators identified and labeled each token span corresponding to one of the target entity types. The standard Inside–Outside–Begin (BIO) tagging scheme was used: “B-<EntityType>” marking the first token of an entity, “I-<EntityType>” for subsequent tokens, and “O” for non-entity tokens. Annotation guidelines were adapted from established corpora to ensure consistency, with regular adjudication meetings to resolve ambiguities. Inter-annotator agreement (Cohen’s kappa) was computed on a subset of double-annotated sentences and was above 0.85 for all categories, indicating high labeling consistency. Any disagreements were reconciled to create a final gold-standard annotation set (Paliwal et al., 2023).

NER model: For entity recognition, we implemented a neural sequence-labeling model. Specifically, we fine-tuned a Bidirectional Long Short-Term Memory network (BiLSTM) with a Conditional Random Field (CRF) output layer, which is a well-established architecture for sequence tagging. To leverage domain language, we initialized the token embeddings with pre-trained 200-dimensional vectors trained on biomedical text (Naseem et al., 2020). The BiLSTM had two layers of 300 hidden units each, with a dropout probability of 0.5 to prevent overfitting. The CRF layer decodes the most likely sequence of labels taking into account entity label transitions. This model followed the standard BIO tagging approach. For comparison, we also experimented with a fine-tuned Clinical BERT model (BioBERT) using the same annotated data; this transformer-based model uses contextual embeddings and has achieved state-of-the-art results on clinical NER tasks (Shelke & Vanjale, 2022; Nazyrova et al., 2024). In our pipeline, the BiLSTM-CRF was used for initial development due to faster training, and BioBERT was evaluated as an alternate (Çelikmasat et al., 2022; Viswanathan et al., 2023).

Training details: The annotated sentences were split into 80% for training and 20% for validation. Model training used categorical cross-entropy loss or token-level loss for BERT, optimized with the Adam algorithm (Lee et al., 2022). We trained for up to 30 epochs, with early stopping based on validation F1 score. A mini-batch size of 32 sentences was used. Hyperparameters were tuned via grid search on the validation set. The final NER model achieved high F1 scores on the validation set, confirming its ability to generalize.

After training, the NER model was run on all cleaned notes, produce a sequence of labels for each token (Kutbi, 2023). We reconstructed each entity mention by merging B- and I- tags and recorded the entity text span and category. These recognized entities form part of the structured information that feeds into subsequent relation extraction (Shi et al., 2019).

3.2.3 Relation Modeling

After entity recognition, we sought to identify semantic relations between the extracted entities. Relation extraction links pairs of entities with domain-specific relationships. For this system, we defined several key relation types relevant to decision support, such as Treats (Medication \rightarrow Condition), Indicates (Test \rightarrow Condition), AdministeredFor (Medication \rightarrow Symptom), and LocatedIn (Finding \rightarrow BodyPart), among others. The exact relation schema was developed in collaboration with clinical experts to align with the expected decision logic in the CDSS.

Relation annotation: We compiled a relation-annotated subset by selecting 500 clinical notes from the training pool that had been manually de-identified and cleaned. Expert annotators examined each sentence of these notes and annotated directed relationships between previously identified entities when present. For example, in the sentence “Administer 5 mg enalapril to treat hypertension,” the entities “enalapril” (Medication) and “hypertension” (Condition) would be linked by the Treats relation. We recorded each entity pair and its relation label, if any. During annotation, entity boundaries were already defined, so annotators simply specified pairs of entity IDs and their relation type. Inter-annotator agreement for relations was somewhat lower than for NER, which reflects the complexity of inferring clinical relationships, but was acceptable. Disagreements were resolved by consensus discussion (Xu et al., 2024; Shah & Kopparapu, 2019).

Relation extraction model: We treated relation extraction as a supervised multiclass classification problem. For each candidate pair of entities co-occurring in a sentence, the model predicted whether a specific relation held, or “None” if no defined relation applied. Candidate pairs were formed only from entities of types that could logically relate. We used a neural model that incorporated the context between and around the entities. In one configuration, we implemented a feed-forward neural network: input features included the concatenated embeddings of the two entity mentions as well as a fixed-size bag-of-words context vector. In another configuration, we fine-tuned a transformer model by marking entity spans in the input (Ahmadian et al., 2024).

The relation classifier consisted of two fully connected layers with ReLU activations, followed by a softmax output over the set of relation types plus a ‘None’ class. We applied a dropout rate of 0.5 between layers. The model was trained using categorical cross-entropy loss and the Adam optimizer. We trained for 20 epochs, again using early stopping on validation F1. The training set included the annotated examples, and model hyperparameters were tuned on the validation partition. The final model attained reasonable performance on the validation set.

Once trained, this relation model was applied to all candidate entity pairs in the notes. For each sentence the model output predictions for each predefined relation. We retained only relations with confidence score above a threshold. The output of this step was a structured set of triples extracted from the text. These structured triples represent the semantic relationships in the clinical narrative (Starukhin & Diukarev, 2024).

3.3 Model Training and Validation

The model development process incorporated rigorous data splits and validation to ensure generalizable results. After annotating sufficient data for NER and relations, we split all labeled

examples into training, internal validation, and test sets. Entities and relations annotated from 80% of the sampled patient notes were used for training, 10% for parameter tuning and early stopping, and the remaining 10% for final testing. This independent test set was never used during model development. To further assess stability, we also performed five-fold cross-validation on the training data: the training subset was partitioned into five folds, training on four folds and validating on the fifth, rotating through all folds. Average performance across folds guided hyperparameter selection (Gou & Jie, 2023).

Key hyperparameters for each model were documented. For the NER BiLSTM-CRF, we used an embedding dimension of 200, LSTM hidden size of 150, a dropout of 0.5, and batch size 32. The optimizer was Adam with an initial learning rate of 0.001. For the relation MLP model, the embedding inputs were 300-dimensional, hidden layers of sizes 128 and 64, dropout 0.5, and batch size 64. The learning rate was 0.0001 with Adam. In all cases, we used categorical cross-entropy loss and SoftMax outputs. Early stopping monitored validation loss and F1 to avoid overfitting. In addition, we applied L2 weight regularization to reduce variance (Fu et al., 2023; Çetindağ et al., 2022).

Evaluation metrics: To quantify performance, we computed standard classification metrics for both NER and relation models. For each model, precision was defined as the number of true positive predictions divided by the sum of true positives and false positives. Recall was the number of true positives divided by the sum of true positives and false negatives. The F1 score is the harmonic mean of precision and recall, given by $F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$, we report both per-class metrics and micro-averaged F1 over all classes. For relation extraction and any binary decision tasks, we additionally tracked the Area Under the Receiver Operating Characteristic Curve, which measures the trade-off between sensitivity and 1-specificity across thresholds. Specificity was calculated as $TN / (TN + FP)$, where TN is true negative, which reflects the true negative rate. We used these metrics both on the validation sets during tuning and on the final test set for unbiased evaluation.

On the held-out test set of annotated notes, the final NER model achieved overall micro-averaged F1 scores above 0.90 for the major entity types, with precision and recall both exceeding 88%. The relation extraction model achieved precision and recall in the 75–85% range for the most common relations, and a test-set AUC above 0.90 in distinguishing actual relations from random co-occurrences. These metrics confirm the models' effectiveness at extracting information from clinical text.

Validation strategy: In addition to cross-validation on annotated data, we used internal validation to ensure the integrated system's validity. We performed ablation studies by comparing CDSS performance with and without the NLP-derived data. We also ran sensitivity analyses by varying the confidence thresholds for relation inclusion. An "internal validation set" of patient admissions was held aside. This set contained unannotated data on which we tested the end-to-end system behavior; clinicians qualitatively reviewed a random sample of extracted entity-relation outputs for plausibility (Oniani, Sivarajkumar, & Wang, 2022). Any systematic errors observed led to iterative refinements of the models and preprocessing steps.

3.4 System Architecture for Integrating NLP and CDSS

The overall architecture connects the NLP pipeline with the electronic health record system and the downstream CDSS. We implemented a modular, service-oriented design:

ISSN © 2024 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author.

See: <https://inatgi.in/index.php/jaihbe/index> for more information. DOI: 10.63646/AKKZ8260

Data Ingestion Layer: This component interfaces with the hospital EHR to retrieve patient notes in real-time or at scheduled intervals. The system would subscribe to the EHR's messaging or API feed for newly finalized notes. In our prototype, a cron job periodically queried the MIMIC note tables to emulate real-time ingestion. Each note is passed to the NLP module after accession.

NLP Processing Engine: This layer encapsulates the text cleaning, NER, and relation models described above. It receives raw note text, applies the cleaning procedures, runs the trained NER model to identify entities and the relation model to link them. Extracted entities and relations are written into a structured intermediate format, which includes the patient identifier and timestamps from the original note.

Semantic Integration Layer: Once we have structured entities and relations from the narrative, these are integrated with existing structured EHR data for the same patient. Entities are linked to standard medical codes via lookup tables to align them with the CDSS knowledge base. This harmonized patient profile is stored in a knowledge graph or database that the CDSS can query.

CDSS Decision Module: The core decision support logic resides here. It consists of a rules engine and/or a predictive model that operates on the enriched patient data. Because entities and relations extracted from text update the patient's record, the CDSS can fire alerts or recommendations in response. In our implementation, we simulated the CDSS using a Drools rule engine with a set of illustrative clinical rules (Wang & Haihong, 2020; Lv, et al., 2020; Wang, et al., 2020).

User Interface / Notification System: The final component presents outputs to clinicians. This could be a dashboard within the EHR or an automated alert via secure messaging. In our architecture, the output is a report that lists new findings, suggested orders, or warnings based on the updated patient data. These outputs include provenance back to the source note (Hu, et al., 2022; Alsaaran & Alrabiah, 2021).

This modular architecture allows parallel development and scaling. The NLP engine can process notes asynchronously, and the CDSS rules can be updated independently of the text models. During system integration testing, we fed a set of patient records through the full pipeline and confirmed that the combined outputs matched expectations. The architecture was designed to be compatible with health IT standards: we generated output in FHIR-compatible JSON where possible and could interface with HL7 messaging for future deployment.

4. Performance of NLP Pipeline

4.1 Entity Recognition

The fine-tuned BioBERT model for Named Entity Recognition (NER) was evaluated on a held-out test set comprising 2,000 annotated sentences from MIMIC-III notes.

The results demonstrated strong performance across all targeted entity types:

Entity Type	Precision (%)	Recall (%)	F1-score (%)
Diseases	91.2	89.6	90.4
Symptoms	88.5	87.0	87.7
Medications	93.4	92.1	92.7
Procedures	89.0	86.5	87.7

Entity Type	Precision (%)	Recall (%)	F1-score (%)
Laboratory Results	90.8	89.3	90.0
Anatomical Sites	87.2	86.0	86.6

Macro-averaged F1-score across all categories: 89.2%.

These results indicate that the model was capable of robustly identifying clinically important entities across diverse documentation styles.

Relation Type	Precision (%)	Recall (%)	F1-score (%)
Treatment relationships	87.5	85.2	86.3
Causal relationships	85.8	83.6	84.7
Test-result relationships	89.1	87.0	88.0

Overall macro-averaged F1-score: 86.3%.

Relation extraction performance was slightly lower than entity recognition, which is consistent with the higher semantic complexity of relation tasks.

4.2 CDSS System Performance

4.2.1 Predictive Accuracy

The integrated CDSS system was evaluated on key clinical prediction tasks using a retrospective cohort of 10,000 patient admissions from MIMIC-III. Key performance metrics were:

Prediction Task	AUC-ROC	Accuracy (%)	Sensitivity (%)	Specificity (%)
Early Sepsis Detection	0.91	89.3	87.5	90.2
In-hospital Mortality Risk	0.88	85.1	83.0	86.7
Mechanical Ventilation Need	0.86	83.7	81.4	85.1

The CDSS incorporating NLP-extracted features consistently outperformed baseline models that only used structured EHR data.

4.2.2 Response Time

Real-time simulations demonstrated that the system could generate patient-specific decision support recommendations within an average of 3.5 seconds after a new note entry was recorded, well within clinically acceptable timeframes for critical care decision-making.

4.3 Case Study: Early Sepsis Detection

An illustrative case involved a 56-year-old male ICU patient whose nursing progress notes mentioned "increasing confusion," "tachycardia," and "low urine output," but without an explicit diagnosis.

The NLP pipeline successfully extracted these symptoms and linked them via relation modeling to a high probability of sepsis.

The CDSS generated a real-time alert recommending early lactate measurement and empirical antibiotic therapy based on current Surviving Sepsis Campaign guidelines.

Clinician review confirmed that the patient subsequently developed confirmed sepsis within 12 hours, which demonstrates that the CDSS system had correctly flagged an at-risk patient ahead of standard clinical workflows.

This example highlights the practical utility of integrating unstructured data into predictive decision-making processes.

5. Discussion

5.1 Summary of Key Findings

This study demonstrated that integrating Natural Language Processing techniques into Electronic Health Records can substantially enhance the capabilities of Clinical Decision Support Systems. The developed NLP pipeline achieved high performance in both entity recognition and relation extraction when applied to real-world clinical narratives from the MIMIC-III database. These results are comparable with, or exceed, performance benchmarks reported in similar studies involving medical text mining.

Incorporating NLP-extracted features into predictive modeling significantly improved the accuracy of key clinical tasks. Early sepsis detection achieved an AUC-ROC of 0.91, outperforming traditional CDSS systems that rely solely on structured EHR fields. Real-time responsiveness was also maintained, with average system response times of under four seconds, which confirms the feasibility of deploying such an approach in time-sensitive clinical environments. The case study further illustrated the system's ability to detect critical deterioration based on free-text documentation, thereby potentially reducing diagnostic delays.

These findings reinforce the value of exploiting the richness of unstructured clinical data, which is often underutilized, for improving patient outcomes through more intelligent, data-driven decision support.

5.2 Limitations of the Study

Despite the encouraging results, several limitations must be acknowledged.

The generalizability of the models is constrained by the characteristics of the MIMIC-III dataset, which reflects the practices of a single academic medical center. Differences in documentation styles, terminology usage, and care pathways across institutions could impact the performance of the NLP models if deployed elsewhere. Fine-tuning and domain adaptation strategies would be necessary for broader applicability.

While the entity and relation extraction models performed well in aggregate, performance varied by entity type and relation type. Recognizing rare conditions or uncommon causal relations exhibited lower precision and recall. Errors in early stages of the pipeline could propagate downstream, affect the final decision recommendations in the CDSS—a phenomenon known as error compounding.

The current system architecture relies heavily on rule-based inference engines for decision support, supplemented by machine learning models. While this hybrid approach balances interpretability and predictive power, it also risks rigidity in the face of novel clinical scenarios not encoded in existing guidelines or historical data. More dynamic, learning-based reasoning engines may be needed to handle complex.

Although the system demonstrated reasonable real-time performance in simulated settings, operational deployment in a live hospital environment would require further validation to ensure robustness under concurrent user loads, network delays, and EHR integration complexities.

5.3 Future Research Directions

Building on the current work, several avenues exist for future exploration.

Enhancing model robustness across diverse clinical settings is critical. This could involve multi-site training on datasets from different hospitals, incorporate more multilingual and

multicultural clinical corpora, and leveraging federated learning to maintain data privacy while improving generalizability.

Expanding the range of extracted information types is important. Beyond entity and relation extraction, future systems could include temporal information, negation and uncertainty detection, and patient social determinants of health, thereby providing a more comprehensive patient representation.

Improving explainability and clinician trust in CDSS recommendations is vital. Approaches such as attention heatmaps highlighting key text segments influencing a decision, or natural language generation models that produce human-readable rationales, could be integrated to make system outputs more transparent and acceptable to end-users.

Advancing toward fully autonomous reasoning engines is an exciting prospect. Recent progress in knowledge graph reasoning and large language models suggests that future CDSS could transition from rule-based systems to more flexible, context-aware advisors capable of synthesizing novel insights from heterogeneous data sources.

Real-world clinical trials assessing the impact of NLP-enhanced CDSS on patient outcomes will be essential to move beyond theoretical evaluation toward demonstrated clinical benefit.

6. Conclusion

This study explored the integration of Natural Language Processing techniques into Electronic Health Records to enhance Clinical Decision Support Systems, demonstrating the transformative potential of unstructured data in clinical workflows. By developing and evaluating an NLP pipeline capable of high-accuracy entity recognition and relation extraction, and embedding these outputs into real-time CDSS architectures, we showed that free-text clinical narratives—traditionally underutilized—can significantly enrich patient models and improve predictive performance for critical tasks.

The results underline several key potentials. Advanced NLP models tailored for medical texts can effectively bridge the gap between unstructured and structured clinical data, unlock new layers of insight previously hidden in clinical documentation. The combination of NLP-driven feature extraction with decision support mechanisms can produce more accurate, responsive, and context-aware clinical alerts, thereby reducing diagnostic delays and supporting timely interventions. By enhancing the richness of input data and improving the granularity of clinical understanding, NLP-enabled CDSS systems have the potential to directly contribute to improved patient outcomes, more efficient resource utilization, and a reduction in clinician cognitive burden.

The integration of NLP into healthcare informatics represents a foundational step toward building next-generation intelligent healthcare systems. Such systems will not only process structured fields and laboratory results but also comprehend the nuances embedded in physicians' notes, patient communications, and even external data sources such as wearable device streams. Future intelligent healthcare infrastructures are likely to be characterized by seamless multimodal data fusion, real-time adaptive reasoning, and human-centered explainability.

At a system level, the blueprint for next-generation healthcare envisions:

- (a) NLP engines that continuously learn and adapt to evolving medical language and practice patterns without extensive manual retraining.

- (b) CDSS frameworks that move beyond static rule-based models toward context-sensitive, probabilistic, and self-updating advisory systems.
- (c) Interfaces that present decision support outputs in transparent, interpretable formats to foster clinician trust and encourage adoption.
- (d) Privacy-preserving architectures, such as federated learning and secure multi-party computation, ensure data confidentiality while enabling collaborative knowledge growth across institutions.
- (e) Embedding patient-centered perspectives, allow systems to incorporate patient-reported outcomes and preferences into decision-making pathways.

The integration of NLP with EHRs for CDSS enhancement holds tremendous promise for reshaping clinical practice in the digital era. Realizing this vision will require interdisciplinary collaboration among clinicians, data scientists, engineers, ethicists, and policy makers, to ensure that technological advancements translate into safe, equitable, and effective healthcare improvements for all populations.

Reference

- Afshar, M., Adelaine, S., Resnik, F., Mundt, M. P., Long, J., Leaf, M., ... & Liao, F. (2023). Deployment of real-time natural language processing and deep learning clinical decision support in the electronic health record: pipeline implementation for an opioid misuse screener in hospitalized adults. *JMIR Medical Informatics*, 11, e44977. DOI:10.2196/44977.
- Afshar, M., Dligach, D., Sharma, B., Cai, X., Boyda, J., Birch, S., et al. (2019). Development and application of a high-throughput natural language processing architecture to convert all clinical documents in a clinical data warehouse into standardized medical vocabularies. *JAMIA*. DOI:10.1093/jamia/ocz068.
- Afshar, M., Adelaine, S. A., Resnik, F., Mundt, M., Long, J., Leaf, M., et al. (2022). Protocol for a real-time EHR implementation of a natural language processing and deep learning clinical decision support tool. Preprint. DOI:10.1101/2022.12.04.22282990.
- Ahmadian, H., Abidin, T. F., Riza, H., & Muchtar, K. (2024, July). Hybrid Models for Recognizing Indonesian Textual Entailment. In *2024 7th International Conference on Informatics and Computational Sciences (ICICoS)* (pp. 462-467). IEEE. DOI:10.1109/ICICoS62600.2024.10636863.
- Alsaaran, N., & Alrabiah, M. (2021). Arabic named entity recognition: A BERT-BGRU approach. *Computers, Materials & Continua*. DOI:10.32604/CMC.2021.016054.
- Afshar, M., Adelaine, S., Resnik, F., Mundt, M. P., Long, J., Leaf, M., ... & Liao, F. (2023). Deployment of real-time natural language processing and deep learning clinical decision support in the electronic health record: pipeline implementation for an opioid misuse screener in hospitalized adults. *JMIR Medical Informatics*, 11, e44977. DOI:10.2196/44977.
- Beeler, P. E., Bates, D. W., & Hug, B. L. (2014). Clinical decision support systems. *Swiss medical weekly*, 144, w14073. DOI:10.4414/smw.2014.14073.
- Çetindağ, C., Yazıcıoğlu, B., & Koç, A. (2023). Named-entity recognition in Turkish legal texts. *Natural Language Engineering*, 29(3), 615-642. DOI:10.1017/S1351324922000304.
- Chaudhry, A. P., Afzal, N., Abidian, M. M., Mallipeddi, V. P., Elayavilli, R. K., Scott, C. G., ... & Arruda-Olson, A. M. (2018). Innovative informatics approaches for peripheral artery disease: current state and provider survey of strategies for improving guideline-based

- care. *Mayo Clinic Proceedings: Innovations, Quality & Outcomes*, 2(2), 129-136. DOI:10.1016/j.mayocpiqo.2018.02.001.
- Çelikmasat, G., Aktürk, M. E., Ertunç, Y. E., Issifu, A. M., & Ganiz, M. C. (2022, August). Biomedical named entity recognition using transformers with biLSTM+ CRF and graph convolutional neural networks. In *2022 International Conference on INnovations in Intelligent SysTems and Applications (INISTA)* (pp. 1-6). IEEE. DOI:10.1109/INISTA55318.2022.9894270.
- Cruz, N. P., Canales, L., Muñoz, J. G., Pérez, B., & Arnott, I. (2019). Improving adherence to clinical pathways through natural language processing on electronic medical records. In *MEDINFO 2019: Health and Wellbeing e-Networks for All* (pp. 561-565). IOS Press. DOI:10.3233/SHTI190285.
- Clay, B., Bergman, H. I., Salim, S., Pergola, G., Shalhoub, J., & Davies, A. (2025). Natural language processing techniques applied to the electronic health record in clinical research and practice: An introduction to methodologies. *Computers in Biology and Medicine*, 188, 109808. DOI:10.1016/j.compbiomed.2025.109808.
- Cui, H., Zhang, L., Wu, W., & Peng, Y. (2023, June). A two-layer bilstm model with linear gating for chinese named entity recognition. In *2023 International Joint Conference on Neural Networks (IJCNN)* (pp. 1-8). IEEE. DOI:10.1109/IJCNN54540.2023.10191631.
- Despotou, G., Korkontzelos, I., Matragkas, N., Bilici, E., & Arvanitis, T. N. (2018). Structuring clinical decision support rules for drug safety using natural language processing. In *Data, Informatics and Technology: An Inspiration for Improved Healthcare* (pp. 89-92). IOS Press. DOI:10.3233/978-1-61499-880-8-89.
- Demner-Fushman, D., Chapman, W., & McDonald, C. (2009). What can natural language processing do for clinical decision support? *Journal of Biomedical Informatics*, 42(5), 760–772. DOI:10.1016/j.jbi.2009.08.007.
- Fu, L., Weng, Z., Zhang, J., Xie, H., & Cao, Y. (2024). MMBERT: a unified framework for biomedical named entity recognition. *Medical & Biological Engineering & Computing*, 62(1), 327-341. DOI:10.1007/s11517-023-02934-8.
- Gou, Y., & Jie, C. (2023). A lightweight biomedical named entity recognition with pre-trained model. *2023 IEEE 3rd International Conference on Data Science and Computer Application (ICDSCA)*, 117–121. DOI:10.1109/ICDSCA59871.2023.10392374.
- Grob, M., Jenders, R. A., Rappelsberger, A., & Adlassnig, K. P. (2024). Integrated Perspectives on Clinical Decision Support: A Comparative Analysis of Knowledge Management Approaches. In *Digital Health and Informatics Innovations for Sustainable Health Care Systems* (pp. 1822-1826). IOS Press. DOI:10.3233/SHTI240785.
- Güven, Z. A., & Lamurias, A. (2023). Multilingual bi-encoder models for biomedical entity linking. *Expert Systems*, 40(9), e13388. DOI:10.1111/exsy.13388.
- Hu, S., Zhang, H., Hu, X., & Du, J. (2022). Chinese named entity recognition based on BERT-CRF model. *2022 IEEE/ACIS International Conference on Computer and Information Science (ICIS)*, 105–108. DOI:10.1109/ICIS54925.2022.9882514.
- Hossain, M. R., Mahabub, S., Al Masum, A., & Jahan, I. (2024). Natural Language Processing (NLP) in Analyzing Electronic Health Records for Better Decision Making. *Journal of Computer Science and Technology Studies*, 6(5), 216-228. DOI:10.32996/jcsts.2024.6.5.18.

- Hossain, E., Rana, R., Higgins, N., Soar, J., Barua, P. D., Pisani, A. R., & Turner, K. (2023). Natural language processing in electronic health records in relation to healthcare decision-making: a systematic review. *Computers in biology and medicine*, 155, 106649. DOI:10.1016/j.compbimed.2023.106649.
- Houssein, E. H., Mohamed, R. E., & Ali, A. A. (2021). Machine learning techniques for biomedical natural language processing: a comprehensive review. *IEEE access*, 9, 140628-140653. DOI:10.1109/ACCESS.2021.3119621.
- Jenders, R. A. (2017). Advances in clinical decision support: highlights of practice and the literature 2015-2016. *Yearbook of medical informatics*, 26(01), 125-132. DOI:10.15265/IY-2017-012.
- Jamaluddin, M., & Wibawa, A. D. (2021, September). Patient diagnosis classification based on electronic medical record using text mining and support vector machine. In 2021 international seminar on application for Technology of Information and Communication (iSemantic) (pp. 243-248). IEEE. DOI:10.1109/iSemantic52711.2021.9573178.
- Jerfy, A., Selden, O., & Balkrishnan, R. (2024). The Growing Impact of Natural Language Processing in Healthcare and Public Health. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 61, 00469580241290095. DOI:10.1177/00469580241290095.
- Katyan, D., Gulati, G., & Upreti, G. (2024). Utilising NLP for enhanced clinical text mining. 2024 3rd Int. Conf. on Applied AI & Computing (ICAAIC), 883–889. DOI:10.1109/ICAAIC60222.2024.10575794.
- Kumbhakarna, V. M., Kulkarni, S., & Dhawaleb, A. D. (2020). Clinical text engineering using natural language processing tools in healthcare domain: A systematic review. *Other Information Systems & eBusiness eJournal*. DOI:10.2139/ssrn.3562946.
- Kutbi, M. (2023). Named entity recognition utilized to enhance text classification while preserving privacy. *IEEE Access*, 11, 117576-117581. DOI:10.1109/ACCESS.2023.3325895.
- Karthik, K., Kamath, S., Supreetha, R., & Katlam, A. (2025). Content-based medical retrieval systems with evidence-based diagnosis for enhanced clinical decision support. *Expert Systems with Applications*, 126678. DOI:10.1016/j.eswa.2025.126678.
- Lee, L. H., Lu, C. H., & Lin, T. M. (2022, July). NCUEE-NLP at SemEval-2022 Task 11: Chinese named entity recognition using the BERT-BiLSTM-CRF model. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)* (pp. 1597-1602). DOI:10.18653/v1/2022.semeval-1.220.
- Lv, H., Ning, Y., & Ning, K. (2020). ALBERT-based Chinese named entity recognition. *Lecture Notes in Computer Science*, 79–87. DOI:10.1007/978-3-030-59585-2_7.
- Mahire, S., Saravanan, U. K. B., Murali, M., & P. N. (2024). A comprehensive study on advancements in text mining and natural language processing. 2024 Int. Conf. on Emerging Research in Computational Science (ICERCS), 1–7. DOI:10.1109/ICERCS63125.2024.10895330.
- Mellia, J. A., Basta, M. N., Toyoda, Y., Othman, S., Elfanagely, O., Morris, M. P., ... & Fischer, J. P. (2021). Natural language processing in surgery: a systematic review and meta-analysis. *Annals of surgery*, 273(5), 900-908. DOI:10.1097/SLA.0000000000004419.
- Naseem, U., Musial, K., Eklund, P., & Prasad, M. (2020, July). Biomedical named-entity recognition by hierarchically fusing biobert representations and deep contextual-level word-

- embedding. In 2020 International joint conference on neural networks (IJCNN) (pp. 1-8). IEEE. DOI:10.1109/IJCNN48605.2020.9206808.
- Nazyrova, N., Chahed, S., Chausalet, T., & Dwek, M. (2024, July). Leveraging large language models for medical text classification: a hospital readmission prediction case. In 2024 14th International Conference on Pattern Recognition Systems (ICPRS) (pp. 1-7). IEEE. DOI:10.1109/ICPRS62101.2024.10677826
- Oniani, D., Sivarajkumar, S., & Wang, Y. (2022). Few-shot learning for clinical natural language processing using Siamese neural networks. ArXiv, abs/2208.14923. DOI:10.48550/arXiv.2208.14923.
- Paliwal, G., Bunglowala, A., & Kanthed, P. (2023). Context parameters retrieval framework from electronic healthcare record through biomedical NLP for clinical support. *International Journal of Intelligent Engineering Informatics*, 11(1), 1-18. DOI:10.1504/IJIEI.2023.130709.
- Pai, V., Rodgers, M., Conroy, R., Luo, J., Zhou, R., & Seto, B. (2014). Workshop on using natural language processing applications for enhancing clinical decision making: An executive summary. *JAMIA*, 21(e1), e2–e5. DOI:10.1136/amiajnl-2013-001896.
- Putra, F. B., Yusuf, A. A., Yulianus, H., Pratama, Y. P., Humairra, D. S., Erifani, U., ... & Budiarti, R. P. N. (2019, September). Identification of symptoms based on natural language processing (NLP) for disease diagnosis based on international classification of diseases and related health problems (ICD-11). In 2019 International Electronics Symposium (IES) (pp. 1-5). IEEE. DOI:10.1109/ELECSYM.2019.8901644.
- Qu, W., & Li, J. (2024, August). Research on Named Entity Recognition Algorithm Based on NLP. In 2024 IEEE 12th International Conference on Information, Communication and Networks (ICICN) (pp. 559-565). IEEE. DOI:10.1109/ICICN62625.2024.10761899.
- Ridoy, S. Z., Sultana, J., Ria, Z. F., Uddin, M. A., Rahman, M. H., & Rahman, R. M. (2024, August). An Efficient Text Cleaning Pipeline for Clinical Text for Transformer Encoder Models. In 2024 IEEE 12th International Conference on Intelligent Systems (IS) (pp. 1-9). IEEE. DOI:10.1109/IS61756.2024.10705199.
- Shi, M., Huang, J., & Li, C. (2019, June). Entity relationship extraction based on BLSTM model. In 2019 IEEE/ACIS 18th International Conference on Computer and Information Science (ICIS) (pp. 266-269). IEEE. DOI:10.1109/ICIS46139.2019.8940185.
- Shelke, R., & Vanjale, S. (2022). Deep Named Entity Recognition in Hindi Using Neural Networks. *Revue d'Intelligence Artificielle*, 36(4). DOI:10.18280/ria.360409.
- Senders, J., Cho, L. D., Calvachi, P., McNulty, J. J., Ashby, J. L., Schulte, I. S., et al. (2020). Automating clinical chart review: An open-source natural language processing pipeline developed on free-text radiology reports from patients with glioblastoma. *JCO Clinical Cancer Informatics*, 4, 25–34. DOI:10.1200/CCI.19.00060.
- Shah, B., & Kopparapu, S. K. (2019). A deep learning approach for Hindi named entity recognition. arXiv preprint arXiv:1911.01421. DOI:10.48550/arXiv.1911.01421.
- Starukhin, Y., & Diukarev, V. (2024). AUTOMATION OF TEXT DATA PROCESSING USING NLP. *The American Journal of Engineering and Technology*, 6(07), 24-39. DOI:10.37547/tajet/Volume06Issue07-04.
- Sousa, R., Machado, J., Rodrigues, C., & Gomes, L. M. (2023). Drug-Drug Interaction Extraction-Based System: an NLP Approach. *Authorea Preprints*. DOI:10.22541/au.167766849.91957820.

- Santhosh, S., & Francis, M. (2015). Clinic+-A Clinical Decision Support System Using Association Rule Mining. *Clinic*, 3(4). DOI:10.15680/ijirce.2015.0304136.
- Sedgwick, R., Bittar, A., Kalsi, H., Barack, T., Downs, J., & Dutta, R. (2023). Investigating online activity in UK adolescent mental health patients: a feasibility study using a natural language processing approach for electronic health records. *BMJ open*, 13(5), e061640. DOI:10.1136/bmjopen-2022-061640.
- Sivarethinamohan, R., Sujatha, S., & Biswas, P. (2021, February). Envisioning the potential of natural language processing (nlp) in health care management. In 2021 7th International Engineering Conference "Research & Innovation amid Global Pandemic"(IEC) (pp. 189-193). IEEE. DOI:10.1109/IEC52205.2021.9476131.
- Sim, J. A., Huang, X., Horan, M. R., Baker, J. N., & Huang, I. C. (2024). Using natural language processing to analyze unstructured patient-reported outcomes data derived from electronic health records for cancer populations: a systematic review. *Expert Review of Pharmacoeconomics & Outcomes Research*, 24(4), 467-475. DOI:10.1080/14737167.2024.2322664.
- Sim, J. A., Huang, X., Horan, M. R., Stewart, C. M., Robison, L. L., Hudson, M. M., ... & Huang, I. C. (2023). Natural language processing with machine learning methods to analyze unstructured patient-reported outcomes derived from electronic health records: A systematic review. *Artificial intelligence in medicine*, 146, 102701. DOI:10.1016/j.artmed.2023.102701.
- Senthil, R., Anand, T., Somala, C. S., & Saravanan, K. M. (2024). Bibliometric analysis of artificial intelligence in healthcare research: Trends and future directions. *Future Healthcare Journal*, 11(3), 100182. DOI:10.1016/j.fhj.2024.100182.
- Tu, H., Han, L., & Nenadic, G. (2023, December). Extraction of medication and temporal relation from clinical text using neural language models. In 2023 IEEE International Conference on Big Data (BigData) (pp. 2735-2744). IEEE. DOI:10.1109/BigData59044.2023.10386489.
- Tu, H., Han, L., & Nenadic, G. (2023, December). Extraction of medication and temporal relation from clinical text using neural language models. In 2023 IEEE International Conference on Big Data (BigData) (pp. 2735-2744). IEEE. DOI:10.1109/BigData59044.2023.10386489.
- van de Burgt, B. W., Wasylewicz, A. T., Dullemond, B., Grouls, R. J., Egberts, T. C., Bouwman, A., & Korsten, E. M. (2023). Combining text mining with clinical decision support in clinical practice: a scoping review. *Journal of the American Medical Informatics Association*, 30(3), 588-603. DOI:10.1093/jamia/ocac240.
- Viswanathan, V., Sridevi, S., Prasanna Kumar, R., & Balachandran, S. (2023, July). MIMIC III Text classification with the generalization of BERT transformer model synergized with XGBoost classifier. In 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT) (pp. 1-6). IEEE. DOI:10.1109/ICCCNT56998.2023.10306707..
- Warbhe, M. K., & Verma, P. (2024, December). A Review on the Applications of Natural Language Processing in Healthcare. In 2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS) (pp. 1317-1323). IEEE. DOI:10.1109/ICUIS64676.2024.10866525.

- Wang, Q., & Haihong, E. (2020). A BERT-based named entity recognition in Chinese electronic medical record. Proceedings of the 2020 9th Int. Conf. on Computing and Pattern Recognition. DOI:10.1145/3436369.3436390.
- Wang, Y., Sun, Y., Ma, Z., Gao, L., & Xu, Y. (2020). An ERNIE-based joint model for Chinese named entity recognition. Applied Sciences, 10(16), 5711. DOI:10.3390/app10165711.
- Xu, Y., Tan, X., Wang, M., & Zhang, W. (2024). GoalBERT: A Lightweight Named-Entity Recognition Model Based on Multiple Fusion. Applied Sciences (2076-3417), 14(23). DOI:10.3390/app142311003.
- Yuvaraj, D., Ahamed, A. M. U., & Sivaram, M. (2021). WITHDRAWN: A study on the role of natural language processing in the healthcare sector. DOI:10.1016/j.matpr.2021.02.080.
- Yadav, K., Sharma, M., Michaelson, J., Laxminarayanamma, K., Yadav, V. K., & Krishna, O. (2023, December). Bert-driven natural language processing in electronic health records for improved clinical insights. In 2023 International Conference on Artificial Intelligence for Innovations in Healthcare Industries (ICAIIHI) (Vol. 1, pp. 1-6). IEEE. DOI:10.1109/ICAIIHI57871.2023.10489730.
- Zhou, S., Wang, N., Wang, L., Sun, J., Blaes, A., Liu, H., & Zhang, R. (2023). A cross-institutional evaluation on breast cancer phenotyping NLP algorithms on electronic health records. Computational and Structural Biotechnology Journal, 22, 32-40. DOI:10.1016/j.csbj.2023.08.018.