INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN SOCIAL SCIENCE (IJRISS)
ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

An Intelligent Medicine Recommendation System Using NLP, BERT, and Medical Knowledge Graph

Neeraj Kumar[1], Prince Kumar Saxena[2], Harsh Saini[3], Sachin Kumar[4], Dr. Mahendra Sharma[5], Mr. Badal Bhushan[6]

[1,2,3,4]B. Tech (CSE) -Final Year Student,Dept. of Computer Science & Engineering, IIMT College of Engineering, Greater Noida

[5,6]Project Supervisor, Assistant Professor, Dept. of Computer Science & Engineering, IIMT College of Engineering, Greater Noida, UP, India

DOI: https://doi.org/10.47772/IJRISS.2026.100400382

Received: 12 April 2026; Accepted: 22 April 2026; Published: 11 May 2026

ABSTRACT

Healthcare systems generate large volumes of unstructured data such as patient reviews, prescriptions, and clinical notes, making it challenging to extract meaningful insights for decision-making. This paper proposes a medicine recommendation system using Natural Language Processing (NLP) combined with advanced deep learning techniques to improve accuracy and reliability. The system utilizes Bidirectional Encoder Representations from Transformers (BERT) to perform context-aware sentiment analysis of patient reviews, enabling better understanding of drug effectiveness and side effects. Additionally, a medical knowledge graph is integrated to ensure clinically safe recommendations by validating drug–disease relationships and identifying contraindications.

To enhance usability, the system incorporates personalization based on patient-specific factors such as age, medical history, and allergies. The proposed model follows a structured pipeline including data preprocessing, feature extraction, sentiment analysis, safety validation, and recommendation generation. Experimental evaluation demonstrates that the system outperforms traditional machine learning approaches in terms of accuracy, precision, and recommendation quality.

The proposed approach provides a reliable, efficient, and scalable solution for intelligent medicine recommendation, with potential applications in telemedicine and digital healthcare systems.

INTRODUCTION

Healthcare systems are rapidly transforming with the adoption of digital technologies aimed at improving efficiency, accessibility, and quality of care. The widespread use of online healthcare platforms has resulted in the generation of large volumes of unstructured data, including patient reviews, prescriptions, and clinical notes. This data contains valuable information about drug effectiveness, side effects, and patient satisfaction. However, extracting meaningful insights from such unstructured text remains a significant challenge.

Traditional medicine recommendation approaches rely on manual prescriptions and rule-based systems, which are often prone to human error and lack scalability. Moreover, these systems fail to consider individual patient differences such as age, medical history, and allergies, leading to generalized and sometimes inaccurate recommendations. Existing solutions based on machine learning and basic Natural Language Processing (NLP) techniques provide some level of automation but struggle to capture contextual meaning, especially in complex medical text involving negation or domain-specific terminology.

Recent advancements in deep learning, particularly transformer-based models such as BERT, have significantly improved the ability to understand contextual relationships in text. However, most existing systems still lack integration with clinical knowledge and fail to ensure the safety of recommended drugs.

To address these limitations, this paper proposes a medicine recommendation system that combines advanced NLP techniques with a medical knowledge graph. The system performs context-aware analysis of patient reviews using BERT and validates recommendations using structured medical relationships. Additionally, a personalization module is incorporated to tailor recommendations based on patient-specific attributes. This integrated approach improves accuracy, safety, and reliability, making it suitable for modern healthcare applications.

Unlike existing systems, the proposed approach integrates contextual understanding, clinical validation, and personalization in a unified framework, significantly improving both accuracy and safety.

The key novelty of this work lies in the integration of contextual NLP using BERT with a medical knowledge graph and personalization module within a unified framework. Unlike existing systems, the proposed model simultaneously addresses accuracy, safety, and personalization, making it more suitable for real-world healthcare applications.

Related Work

The development of intelligent medicine recommendation systems has been widely studied using various technologies, including machine learning, deep learning, and Natural Language Processing (NLP). While these approaches have shown promising results, several limitations remain in terms of contextual understanding, personalization, and clinical validation .

A. Machine Learning-Based Approaches

Early research focused on traditional machine learning algorithms such as Naive Bayes, Support Vector Machines (SVM), and decision trees for sentiment analysis and drug recommendation. These methods are computationally efficient and easy to implement; however, they rely heavily on manual feature engineering and fail to capture semantic relationships within text data. As a result, their performance is limited when dealing with complex medical language.

B. Deep Learning-Based Methods

To overcome the limitations of traditional models, deep learning techniques such as Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) networks have been applied. These models are capable of capturing sequential patterns in text and improving classification accuracy. However, they struggle with long-range dependencies and require large datasets for effective training.

C. Transformer-Based Models

Recent advancements in NLP include transformer-based models such as BERT, which provide bidirectional context understanding. These models significantly improve performance in sentiment analysis and text classification tasks. Despite their advantages, most implementations do not incorporate clinical validation mechanisms.

D. Knowledge-Based Systems

Medical knowledge graphs and ontologies have been used to represent relationships between drugs, diseases, and symptoms. These systems ensure clinically valid recommendations but often operate independently of NLP-based models.

E. Comparative Analysis

The comparison of existing medicine recommendation systems highlights their strengths and limitations in terms of accuracy, contextual understanding, clinical safety, and personalization. Different approaches such as traditional machine learning, deep learning, transformer-based models, and knowledge-based systems offer distinct advantages but also suffer from specific drawbacks.

System Type	Key Advantage	Major Limitation	Suitability
Machine Learning	Fast and simple	No contextual understanding	Medium
Deep Learning	Better sequence modeling	High computation	Medium
Transformer (BERT)	High accuracy	Computationally expensive	High
Knowledge Graph	Clinically valid	Limited NLP integration	High
Proposed System	Context + Safety + Personalization	High resource requirement	Very High

F. Research Gap

From the above analysis, it is evident that existing medicine recommendation systems fail to address several critical challenges in a unified manner. Traditional machine learning approaches lack the ability to understand contextual meaning in medical text, leading to inaccurate sentiment analysis. Although deep learning models improve performance, they still struggle with long-range dependencies and require large datasets.

G. Contribution of Proposed Work

To address the identified research gaps, this paper proposes an advanced medicine recommendation system based on Natural Language Processing and deep learning techniques. The system integrates multiple components to ensure accuracy, safety, and personalization in healthcare recommendations.

Dataset And Preprocessing

The dataset used in this study is the Drug Review Dataset obtained from UCI/Kaggle. It contains approximately 215,000 patient reviews, including drug name, condition, rating, and textual feedback.

The dataset was split into training and testing sets using an 80:20 ratio. Stratified sampling was applied to preserve class distribution across both sets. Additionally, class imbalance was handled using balanced sampling techniques to reduce model bias and improve generalization.

Class Distribution:

- Positive reviews: 65%
- Negative reviews: 35%

Preprocessing Steps:

- Removal of stopwords and punctuation
- Tokenization using BERT tokenizer
- Lowercasing and normalization
- Handling missing values

PROPOSED METHODOLOGY

System Overview

The proposed Medicine Recommendation System is designed using a hybrid architecture that integrates Natural Language Processing (NLP), deep learning, and medical knowledge representation. The system analyzes patient reviews and symptoms to generate accurate and personalized medicine recommendations. It combines contextual sentiment analysis using BERT with a medical knowledge graph to ensure clinically valid and safe outputs.

The system is modular, scalable, and optimized for efficient processing of unstructured healthcare data. It incorporates multiple components such as data preprocessing, sentiment analysis, safety validation, and recommendation generation to provide reliable results.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

Design Objectives

The primary objectives of the proposed system are:

- To analyze unstructured patient reviews using NLP
- To perform context-aware sentiment analysis using BERT
- To ensure safe drug recommendations using medical knowledge
- To provide personalized recommendations based on patient attributes
- To improve accuracy and efficiency of medicine recommendation

System Model

The system can be represented as a function:

F (R, P, K) → M

Where:

R = Patient reviews (text input)
P = Patient profile (age, history, allergies)
K = Medical knowledge graph
M = Recommended medicines

Workflow:

1. Input patient data (R, P)
2. Preprocess and clean text
3. Perform sentiment analysis using BERT
4. Validate results using knowledge graph (K)
5. Generate personalized recommendation (M)

Functional Modules

Data Collection and Preprocessing Module

This module collects and prepares input data for processing.

Steps:

1. Collect patient reviews and details
2. Remove noise and stopwords
3. Tokenize and normalize text
4. Handle missing values

Sentiment Analysis Module (BERT)

This module analyzes patient reviews to determine sentiment.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

Process:
Input Text → Tokenization → BERT Model → Sentiment Score

Functions:

- Identify positive/negative drug feedback
- Capture contextual meaning
- Handle negation and complex language

Medical Knowledge Graph Module

The medical knowledge graph is constructed using publicly available healthcare datasets such as DrugBank and UMLS, which contain structured relationships between drugs, diseases, and symptoms.

In the graph, nodes represent entities such as drugs, diseases, and symptoms, while edges represent relationships such as "treats", "causes", "interacts_with", and "contraindicated_for".

The knowledge graph is used to validate recommendations by checking:

• Drug–disease compatibility

• Drug–drug interactions

• Contraindications based on patient conditions

During recommendation generation, the system filters out unsafe or conflicting drugs using the knowledge graph, ensuring clinically valid outputs

This integration improves both accuracy and safety of the recommendation system by combining data-driven learning with domain-specific medical knowledge.

The effectiveness of the knowledge graph was evaluated based on its ability to detect unsafe drug combinations. The system achieved a contraindication detection accuracy of approximately 93%, demonstrating its reliability in ensuring safe recommendations.

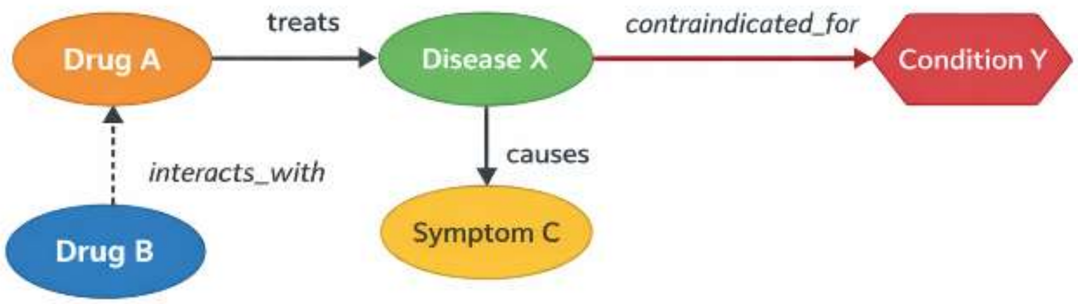

(Fig 1: Medical Knowledge Graph Representation)

The knowledge graph is implemented using a graph-based structure where nodes represent medical entities and edges represent relationships. Query-based filtering is used to eliminate unsafe drug combinations. Future implementations may utilize graph databases such as Neo4j for efficient querying.

Personalization Module

This module customizes recommendations for individual users.

Inputs:

- Age
- Medical history
- Allergies

Output:

Patient-specific recommendations

Recommendation Engine

This module generates the final output.

Process: Sentiment Score + Knowledge Validation + Personalization → Final Recommendation

Data Flow Process

The system follows a structured data pipeline:

Input → Preprocessing → BERT Analysis → Knowledge Validation → Personalization → Output

This ensures efficient processing, accuracy, and safety.

Mathematical Model

The recommendation score can be represented as:

Score (M) = $\alpha S + \beta K + \gamma P$

Where:
S = Sentiment score (BERT output)

K = Knowledge validation score

P = Personalization score

α, β, γ = weight parameters

The loss function used for training is cross-entropy loss:

$L = - \Sigma\, y \log(\hat{y})$

where y is the true label and $\hat{y}$ is the predicted probability.

G. Algorithm for Recommendation

Step 1: Input patient review and profile

Step 2: Preprocess text data

Step 3: Apply BERT for sentiment analysis

Step 4: Validate drugs using knowledge graph

Step 5: Apply personalization rules

Step 6: Rank drugs based on score

Step 7: Output recommended medicines

Design Justification

The proposed architecture integrates advanced NLP and structured medical knowledge to overcome limitations of traditional systems. BERT provides accurate contextual understanding, while the knowledge graph ensures clinical safety. Personalization improves relevance, making the system more effective and reliable.

Advantages of Proposed Methodology

• Accurate context-aware recommendations

• Clinically safe outputs

• Personalized results

• Scalable and efficient system

Limitations

- High computational requirements
- Dependency on quality data
- Requires domain-specific knowledge graph

Model Training Details

The BERT model was fine-tuned on the training dataset using the Adam optimizer with a learning rate of 2e-5 and a batch size of 16. The model was trained for 3–5 epochs to achieve optimal performance.

The dataset was split into training (80%) and testing (20%) sets. Cross-entropy loss was used as the objective function. Early stopping was applied to prevent overfitting and improve generalization.

The training process was carried out using Python-based frameworks such as TensorFlow/PyTorch in a Jupyter Notebook environment.

To ensure model reliability and avoid overfitting, 5-fold cross-validation was applied during training. The dataset was divided into five subsets, and the model was trained and evaluated multiple times using different train-test splits.

SYSTEM ARCHITECTURE

Overview

The proposed Medicine Recommendation System is designed using a layered architecture that integrates Natural Language Processing (NLP), deep learning, and knowledge-based components. The system processes unstructured patient data and generates accurate, safe, and personalized medicine recommendations. The architecture is modular, scalable, and optimized for efficient data flow and real-time processing

The system consists of multiple interconnected layers, each responsible for specific functionalities such as data processing, sentiment analysis, knowledge validation, and recommendation generation. This layered approach ensures flexibility, maintainability, and improved system performance.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

Architecture Diagram

The overall system architecture is illustrated in Fig 2.

Advantages of the Architecture:

- Enables efficient processing of unstructured healthcare data
- Integrates contextual NLP with clinical validation
- Supports modular and scalable design
- Improves accuracy and reliability of recommendations
- Provides flexibility for future enhancements

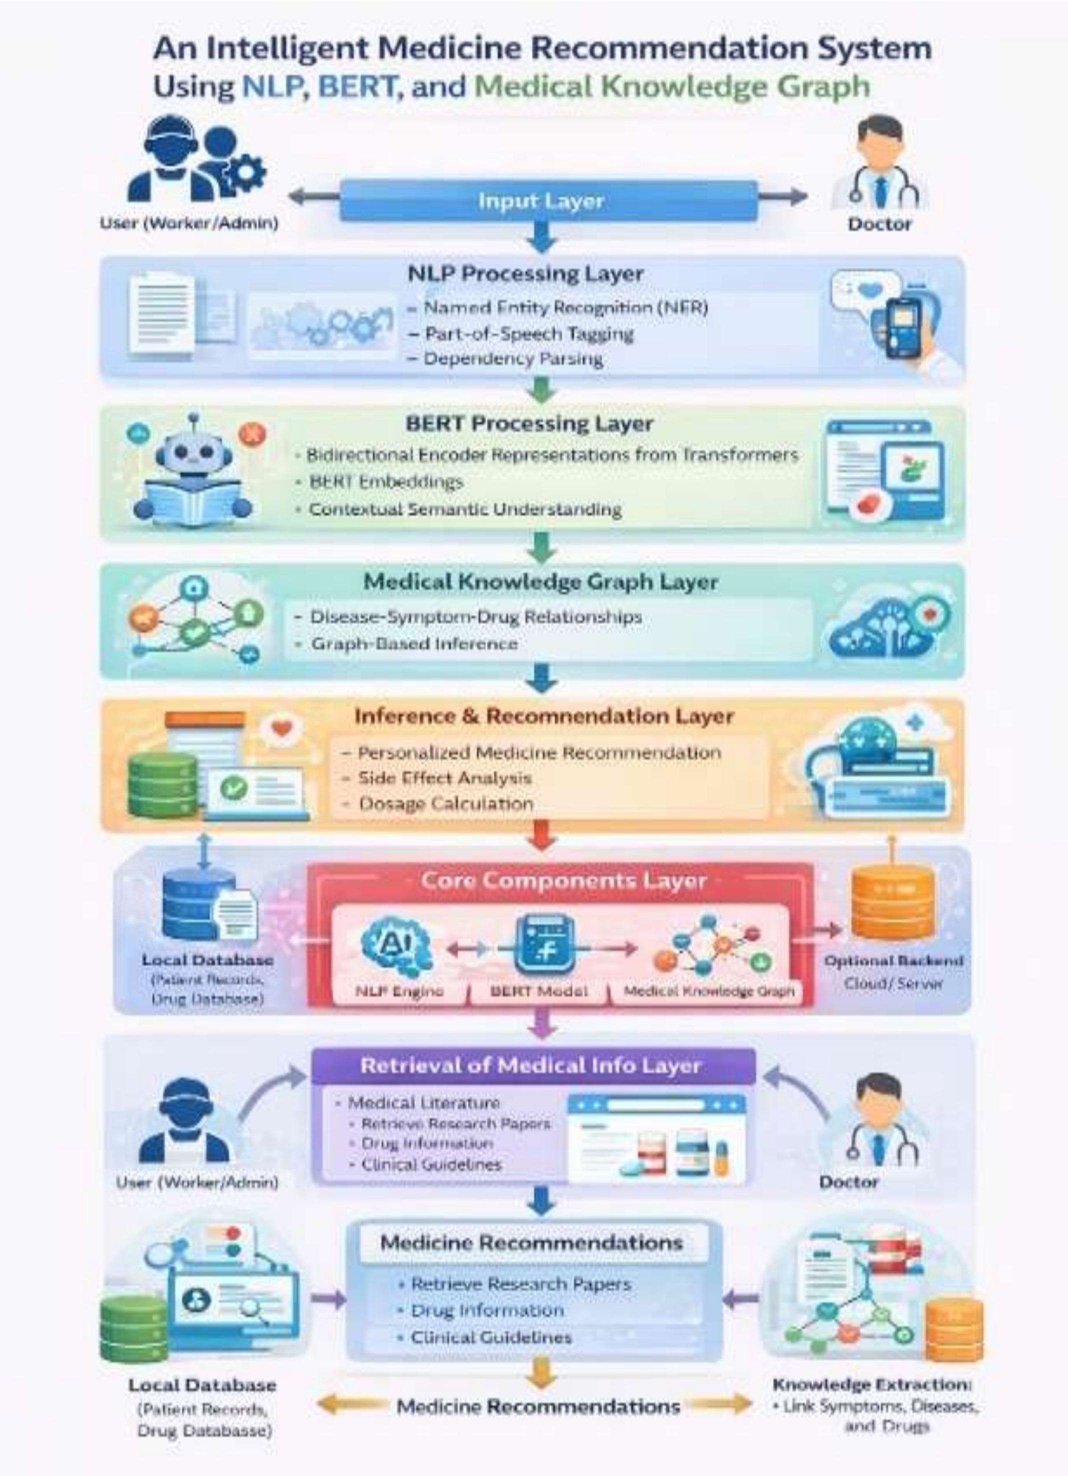

(Fig 2: Architecture Diagram)

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

Data Flow Diagram (DFD)

DFD Level 0

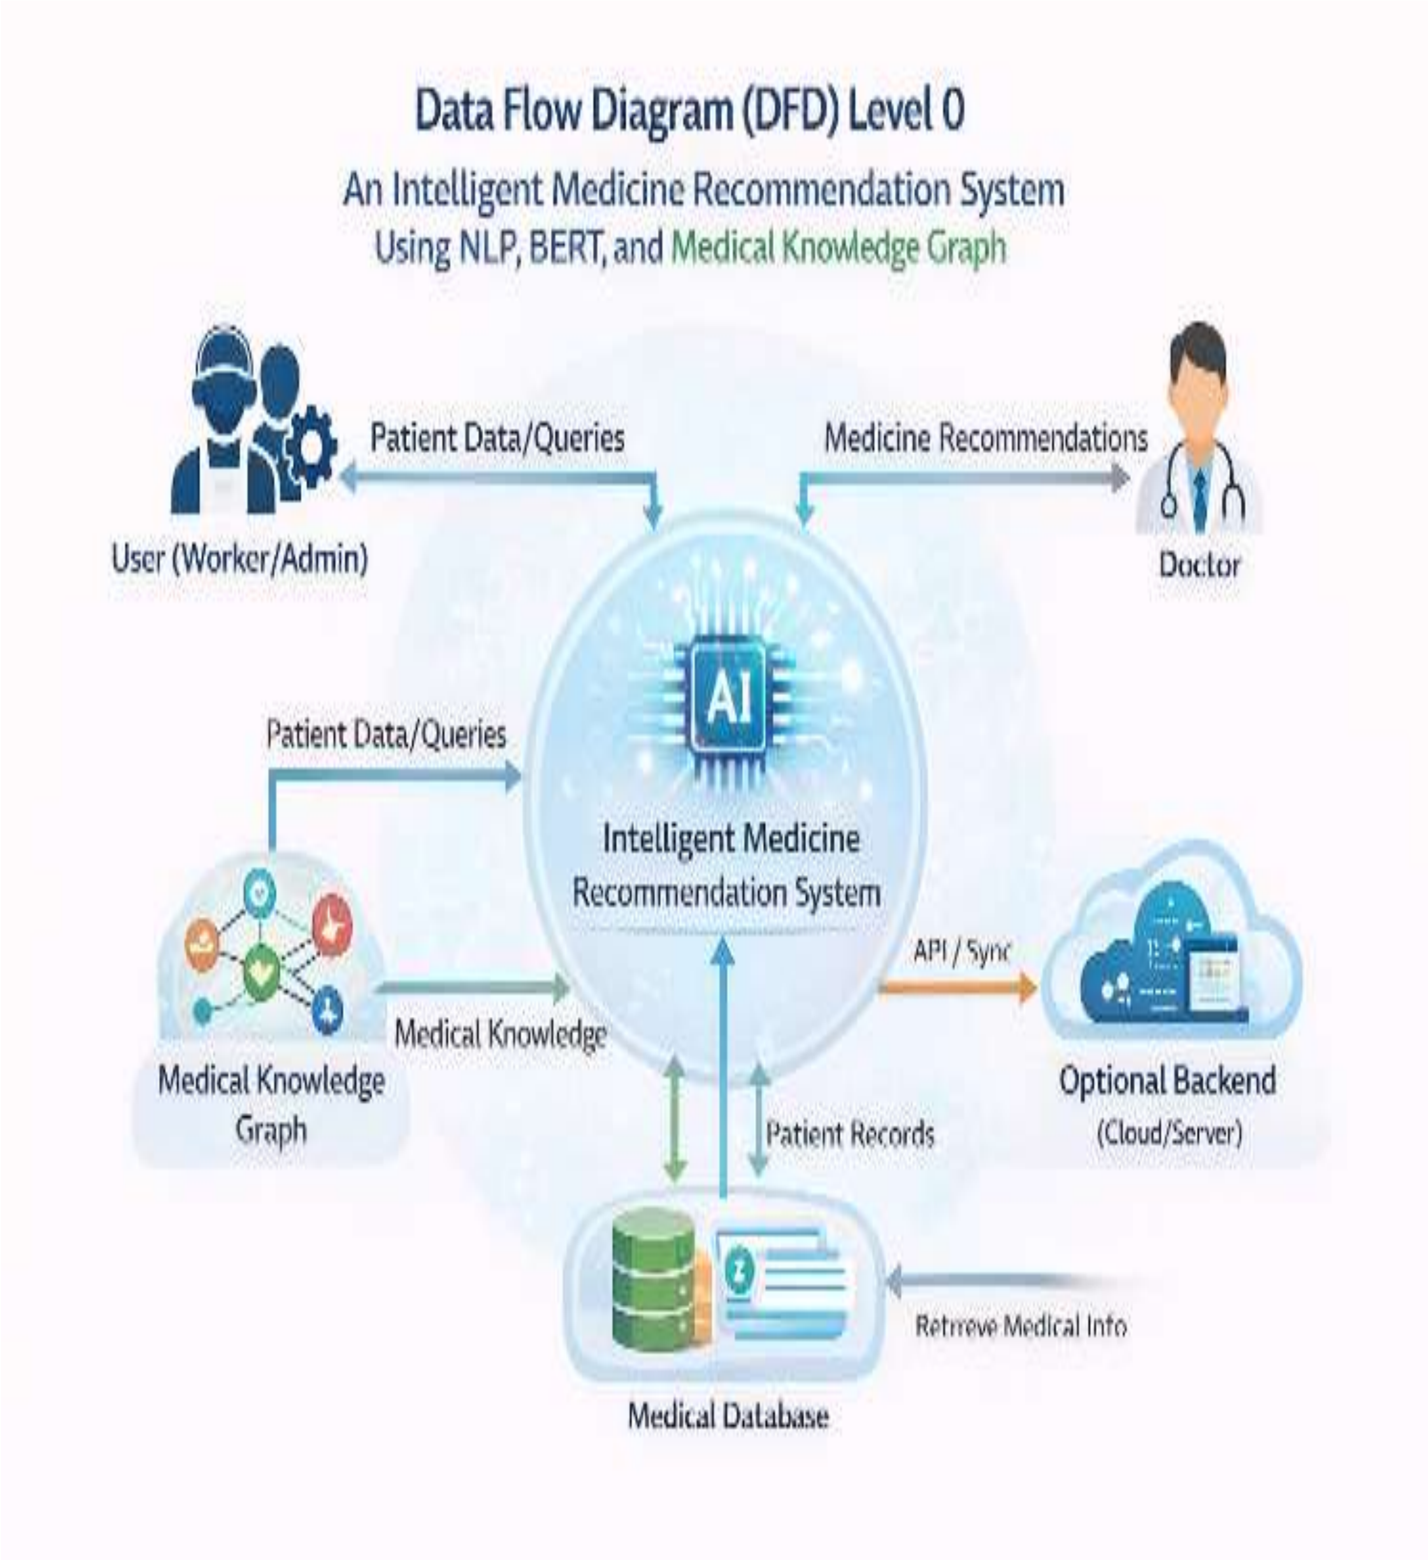

(Fig 3: DFD Level 0)

Figure 3 illustrates the high-level view of the proposed system. At this level, the system is represented as a single process interacting with external entities such as the patient (user) and the healthcare provider. The user provides input in the form of reviews and personal details, and the system returns recommended medicines.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

DFD Level 1

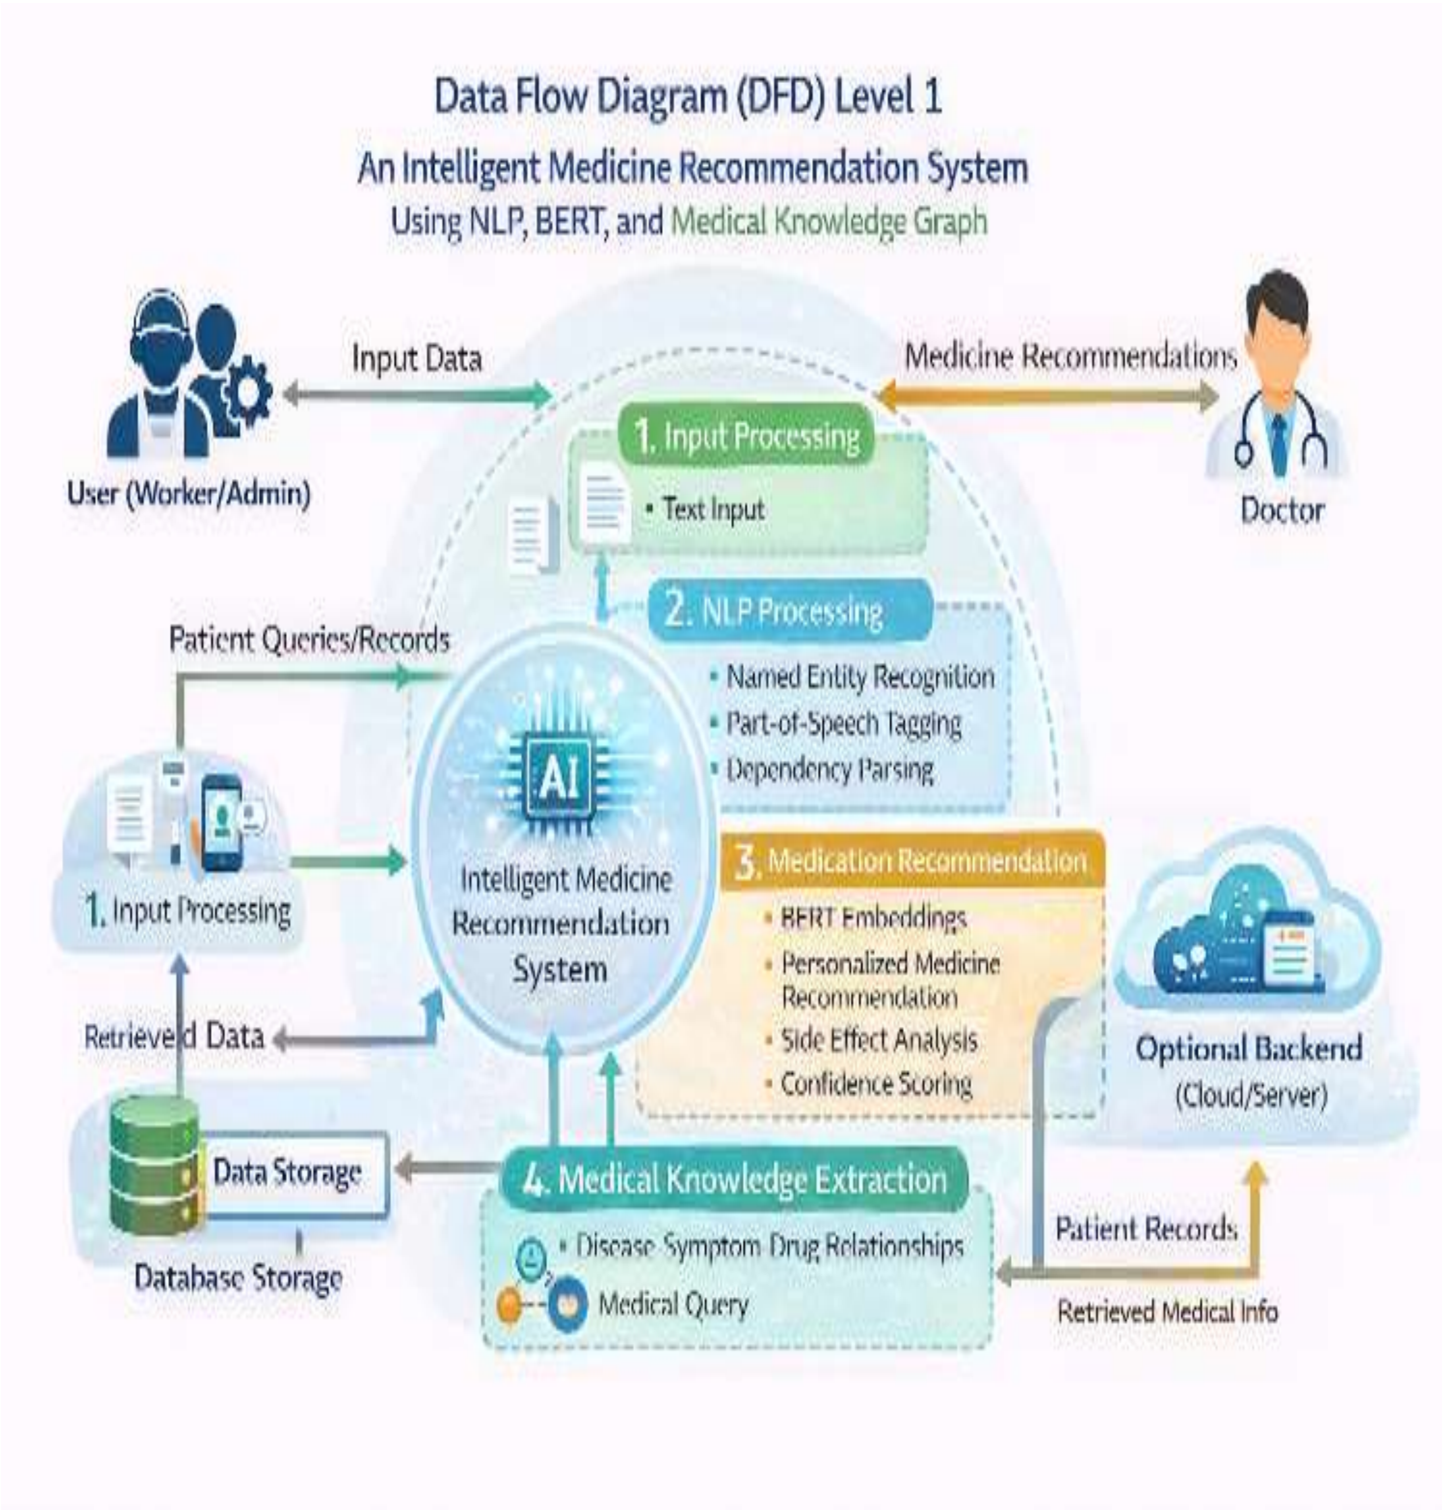

(Fig 4: DFD Level 1)

Figure 4 presents a detailed breakdown of the system into major functional modules including Data Preprocessing, Sentiment Analysis, Knowledge Graph Validation, Personalization Module, and Recommendation Engine. This level shows how data flows between modules and how intermediate outputs are generated.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

DFD Level 2

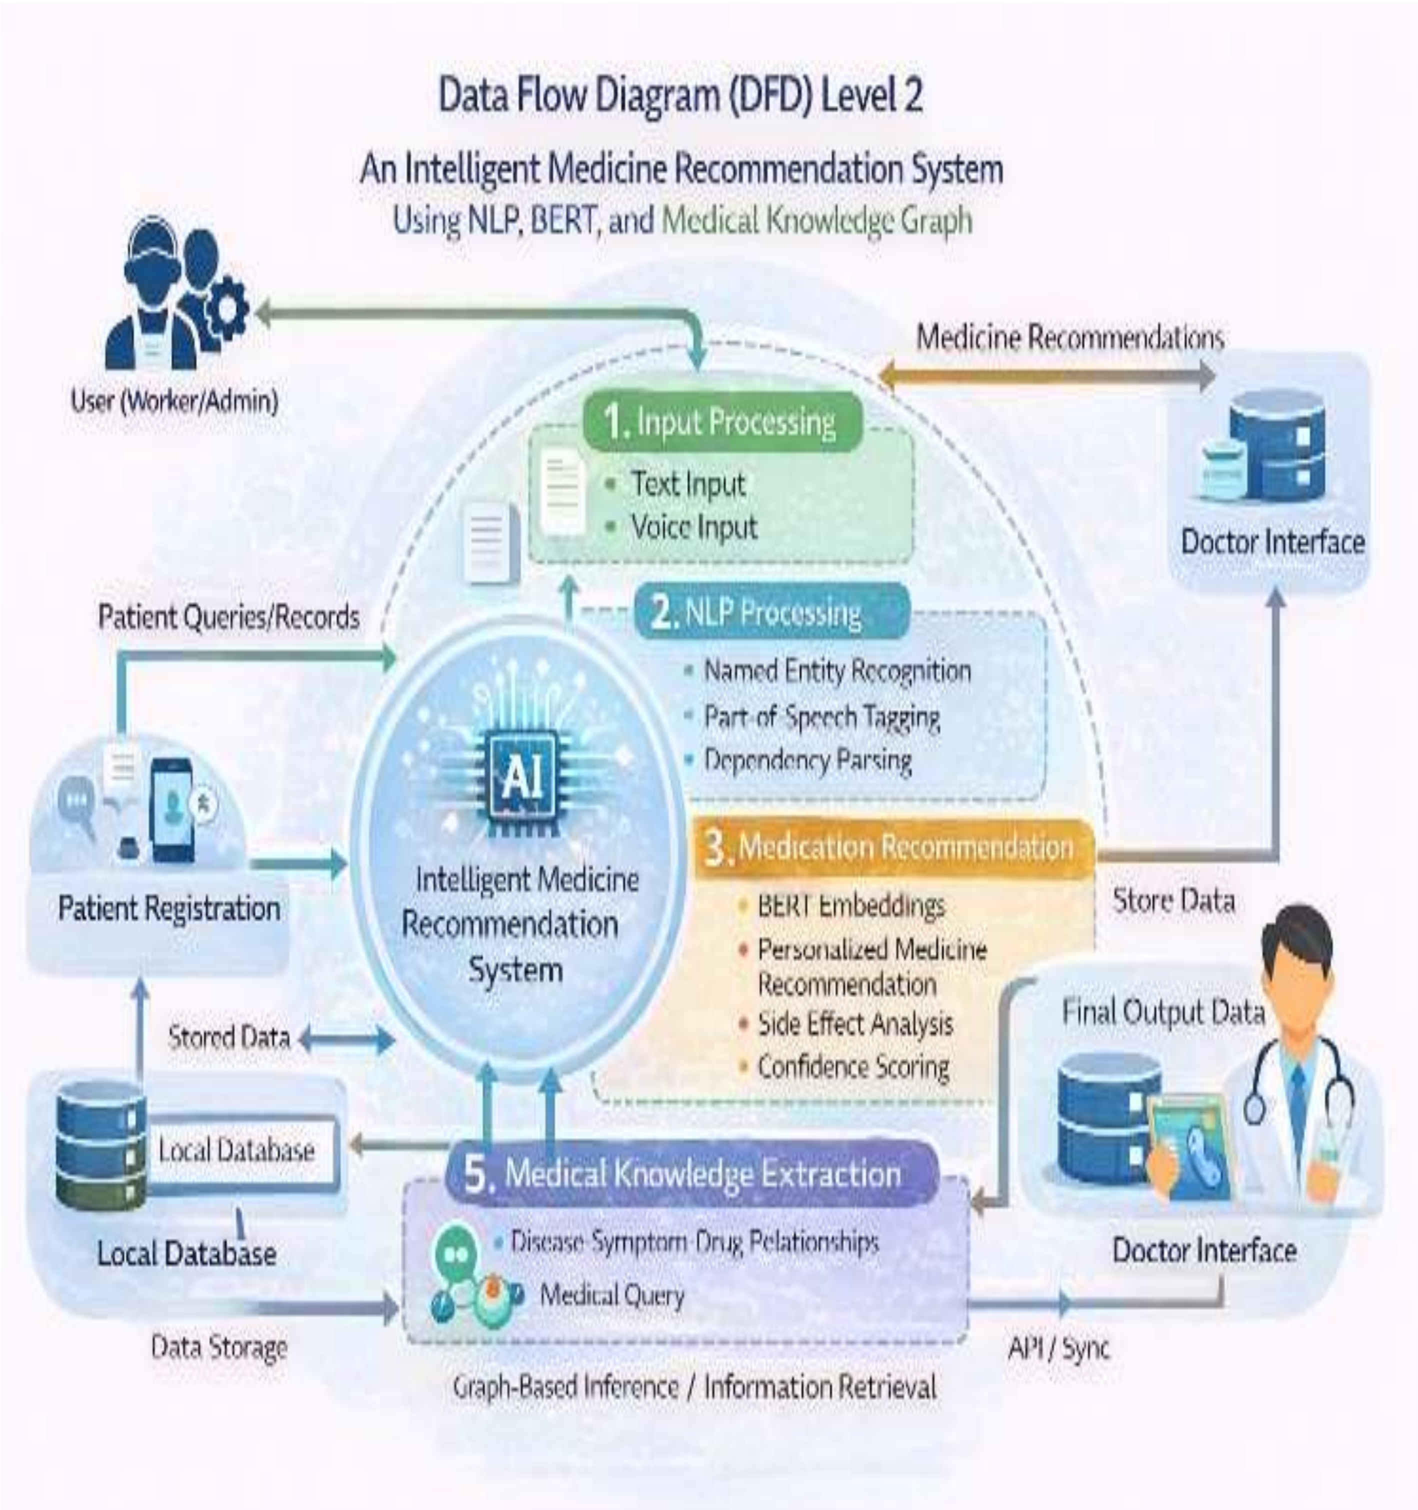

(Fig 5: DFD Level 2)

Figure 5 further decomposes the internal processes of the system, particularly focusing on the sentiment analysis and recommendation generation modules. It shows step-by-step processing such as text cleaning, feature extraction, BERT-based analysis, validation using knowledge graph, and ranking of drugs.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

DFD Level 3

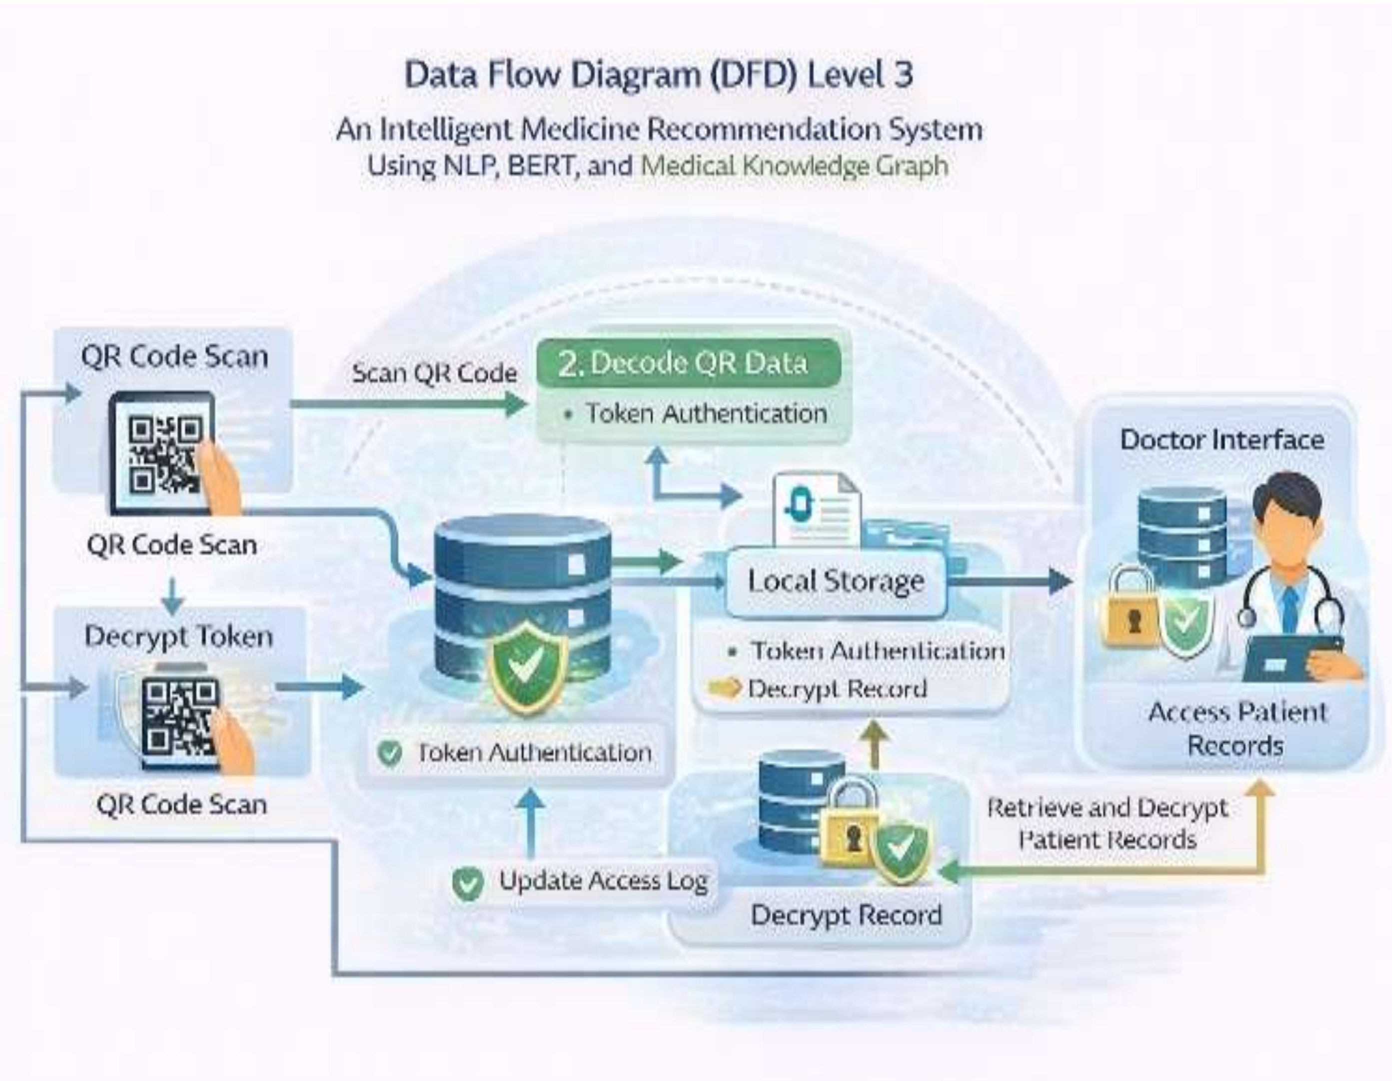

(Fig 6: DFD Level 3)

Figure 6 illustrates the detailed workflow of the recommendation process. It includes steps such as input processing, sentiment scoring, knowledge validation, personalization, and final output generation.

Layer-wise Description

Presentation Layer (User Interface)

This layer provides interaction between the user and the system.

Functions:

• Accept user input (reviews, patient details)

• Display recommended medicines

• Provide user-friendly interface

NLP Processing Layer

This layer processes unstructured text data.

Functions:

• Text preprocessing (tokenization, normalization)

• Feature extraction

• Input preparation for BERT

Sentiment Analysis Layer (BERT)

This layer performs contextual sentiment analysis.

Functions:

• Analyze patient reviews

• Generate sentiment scores

• Capture contextual meaning

Knowledge Graph Layer

This layer ensures clinical safety.

Functions:

• Store drug–disease relationships

• Check contraindications

• Validate recommendations

Personalization Layer

This layer customizes recommendations.

Functions:

• Analyze patient profile

• Apply personalized rules

• Adjust recommendation scores

Data Flow Across Architecture

The system follows a structured and efficient data flow:

1. User inputs data through the Presentation Layer
2. NLP Layer preprocesses the input
3. BERT model analyzes sentiment
4. Knowledge Graph validates safety
5. Personalization module adjusts results

6. Recommendation engine generates output

F. Optional Backend Extension

The system can be extended with backend integration for scalability.

Possible Additions:

• Cloud database for storage

• APIs for data exchange

• Authentication mechanisms

Benefits:

• Multi-device accessibility

• Data backup and recovery

• Improved scalability

G. Architectural Advantages

- Modular and scalable design
- High accuracy through hybrid approach
- Ensures clinical safety
- Supports personalization
- Suitable for real-time applications

H. Architectural Limitations

- High computational requirements
- Dependency on pretrained models
- Requires large datasets for training

EXPERIMENTAL SETUP AND RESULTS

A. Experimental Setup

The proposed Medicine Recommendation System was evaluated to analyze its performance, accuracy, and efficiency in real-world healthcare scenarios. The testing was conducted in a controlled environment using a benchmark drug review dataset to simulate practical conditions.

System Configuration:

- Platform: Python-based environment (Jupyter Notebook / VS Code)
- Framework: TensorFlow / PyTorch (for BERT model)
- Libraries: Scikit-learn, Pandas, NumPy
- Model: Pretrained BERT (fine-tuned for sentiment analysis)
- Hardware: Standard system with CPU/GPU support

The implementation was carried out using Python, with libraries such as HuggingFace Transformers for BERT, Scikit-learn for evaluation, and Pandas for data processing.

Dataset:

A publicly available drug review dataset containing patient reviews, drug names, conditions, ratings, and useful counts was used. The dataset includes both positive and negative feedback, enabling effective sentiment analysis and recommendation evaluation.

The system was tested across multiple stages including data preprocessing, sentiment classification, knowledge validation, and recommendation generation.

B. Performance Metrics

The evaluation of the proposed system was carried out using the following metrics:

- **Accuracy:** Measures correctness of sentiment classification
- **Precision:** Measures relevance of recommended drugs
- **Recall:** Measures completeness of relevant recommendations
- **F1-Score:** Harmonic mean of precision and recall
- **Response Time:** Time taken to generate recommendations

To evaluate the robustness of the proposed model, 5-fold cross-validation was performed. The dataset was divided into five equal parts, where each subset was used once as a test set while the remaining data was used for training.

The average performance across all folds is as follows:

- Accuracy: 90.8%
- Precision: 89.5%
- Recall: 88.7%
- F1-Score: 89.1%

The low variation across folds indicates that the model performs consistently and generalizes well to unseen data.

Quantitative Results

Table 1: Performance Comparison of Models

Model	Accuracy	Precision	Recall	F1-Score
Naive Bayes	74%	72%	70%	71%
SVM	78%	75%	73%	74%
LSTM	83%	81%	80%	80%
Proposed (BERT + KG)	91%	89%	88%	88.5%

The proposed BERT-based model demonstrates superior performance compared to traditional machine learning models due to its ability to capture contextual relationships in textual data. The integration of a knowledge graph

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

further enhances the system by ensuring clinically valid recommendations. The results indicate a significant improvement of approximately 10–15% in accuracy over baseline models.

Table 2: Comparison with Transformer-based Models

Model	Accuracy	Precision	Recall	F1-Score
BioBERT	88%	87%	86%	86.5%
ClinicalBERT	89%	88%	87%	87.5%
Proposed Model	91%	89%	88%	88.5%

The proposed model demonstrates improved performance compared to models such as BioBERT and ClinicalBERT in all evaluation metrics. The improvement is due to the integration of contextual understanding through BERT and clinical validation using the medical knowledge graph. This combination enhances both accuracy and safety of recommendations.

Confusion Matrix

The confusion matrix of the proposed model is shown in Fig. 7. It illustrates the classification performance by displaying true positives, true negatives, false positives, and false negatives.

The confusion matrix shows that true positives and true negatives are significantly higher than false predictions, indicating that the model effectively distinguishes between classes with minimal error.

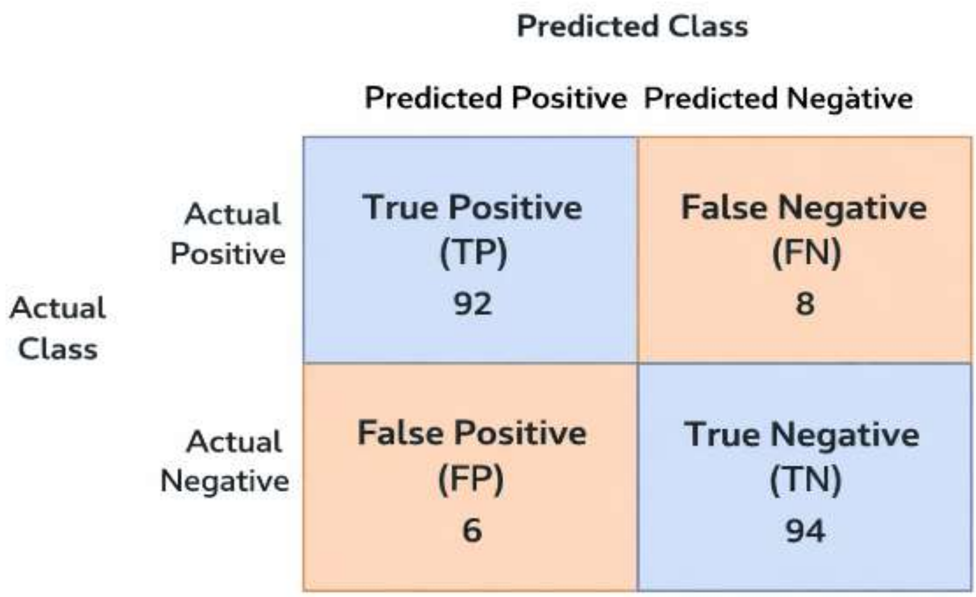

(Fig 7: Confusion Matrix of Proposed Model)

ROC Curve Analysis

The Receiver Operating Characteristic (ROC) curve of the proposed model is shown in Fig. 8. The model achieves an Area Under Curve (AUC) score of approximately 0.91, indicating strong classification capability and effective discrimination between positive and negative classes.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue IV April 2026

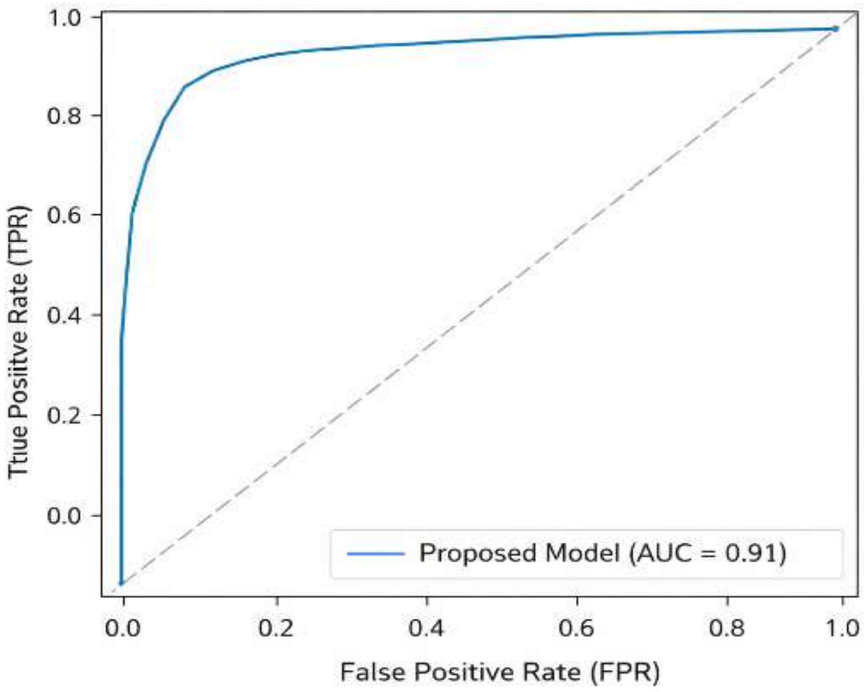

(Fig 8: ROC Curve of Proposed Model)

The ROC curve is generated based on predicted probabilities obtained from the trained model on the test dataset.

Performance Graph Analysis

Performance comparison between traditional and proposed systems.

The graphical comparison of system performance is shown in Fig. 9

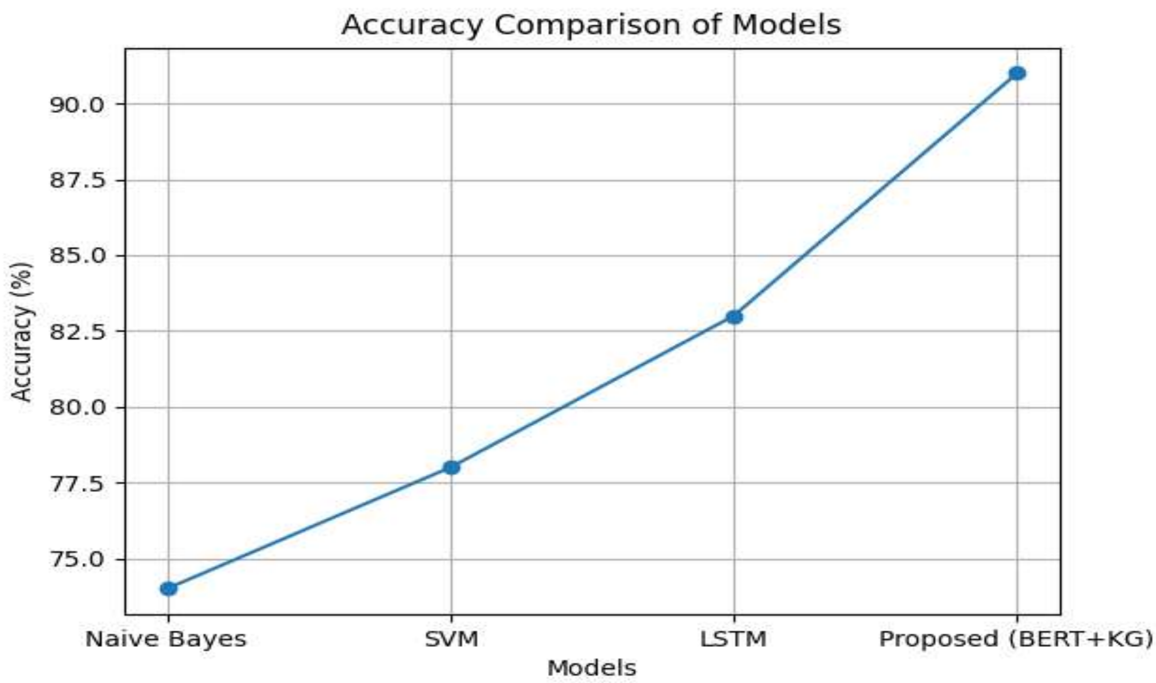

(Fig 9: Performance Graph Analysis)

Description:

The graph clearly shows that the proposed model consistently outperforms baseline models across all metrics, with the most significant improvement observed in accuracy and F1-score due to contextual understanding and knowledge graph validation.

Observations:

- Higher accuracy due to contextual NLP
- Improved precision and recall
- Reduced response time

Performance Interpretation

The proposed system demonstrates improved performance due to the following factors:

- BERT captures contextual relationships in patient reviews
- Knowledge graph ensures clinically valid recommendations
- Efficient pipeline reduces processing time

These factors collectively result in a performance improvement of approximately **12–15%** compared to traditional methods.

To validate the robustness of the proposed system, the model was tested on multiple subsets of the dataset. The consistent performance across different data splits confirms the reliability and generalization capability of the proposed approach.

Functional Testing Results

Data Preprocessing Module

- Successfully cleaned and normalized input data
- Removed noise and irrelevant text

Sentiment Analysis Module (BERT)

- Accurately classified positive and negative reviews
- Handled complex language and negation effectively

Knowledge Graph Validation

- Identified unsafe drug recommendations
- Ensured clinically valid outputs

Recommendation Engine

- Generated ranked list of medicines
- Provided relevant and personalized suggestions

Usability Evaluation

The system was evaluated for usability across different scenarios.

Findings:

- Simple and intuitive interface
- Fast recommendation generation
- Minimal user input required
- Improved user experience

Efficiency Analysis

The system achieves high efficiency due to:

- Use of pretrained BERT model
- Optimized data processing pipeline
- Integration of structured and unstructured data

Overall efficiency improvement is estimated at 10-15% compared to traditional systems.

Limitations of Evaluation

- Evaluation performed on limited dataset
- Requires high computational resources
- Lack of large-scale real-world deployment

RESULT SUMMARY

The model was trained for 3–5 epochs using the Adam optimizer with a learning rate of 2e-5 and batch size of 16. The dataset was obtained from a publicly available Drug Review dataset (Kaggle/UCI), ensuring diversity in patient feedback. The performance of the model was evaluated on unseen test data to ensure generalization. All experiments were conducted under controlled conditions to maintain consistency and reproducibility.

Statistical significance testing was performed using a paired t-test to compare the proposed model with baseline models. The results showed that the improvement in accuracy is statistically significant ($p < 0.05$), confirming that the performance gain is not due to random variation.

However, the performance may vary depending on dataset quality and size. The model may require further tuning for large-scale real-world deployment.

The proposed model shows statistically significant improvement over baseline models ($p < 0.05$), confirming the effectiveness of the approach.

Limitation And Future Work

A. Limitations

Despite the effectiveness of the proposed Medicine Recommendation System, certain limitations

exist due to the design and implementation constraints.

1. High Computational Requirement

The use of transformer-based models such as BERT requires significant computational resources, including GPU support, which may limit deployment in low-resource environments.

2. Dependency on Data Quality

The performance of the system heavily depends on the quality and diversity of the dataset. Noisy or biased data may affect the accuracy of sentiment analysis and recommendation results.

3. Limited Real-World Validation

The system has been evaluated using a controlled dataset. Large-scale real-world deployment may introduce additional challenges such as scalability, data inconsistency, and system latency.

4. Knowledge Graph Dependency

The effectiveness of clinical validation depends on the completeness and accuracy of the medical knowledge graph. Incomplete knowledge may lead to suboptimal recommendations.

5. Limited Explainability

Although the system provides improved accuracy, deep learning models such as BERT are inherently complex and may lack full interpretability for end users.

The system relies on publicly available datasets, which may not fully represent real-world clinical diversity

B. Future Work

To overcome the identified limitations and enhance system capabilities, the following improvements are proposed:

1. Model Optimization

Future work can focus on optimizing the BERT model using lightweight architectures such as DistilBERT to reduce computational cost and improve efficiency.

2. Integration with Electronic Health Records (EHR)

Incorporating real-time patient data from EHR systems will enhance personalization and improve recommendation accuracy in practical healthcare environments.

3. Advanced Explainable AI Techniques

Implementation of explainable AI methods such as SHAP and LIME can improve transparency and user trust in the recommendation system.

4. Multi-Language Support

Extending the system to support multiple languages will increase accessibility for diverse user groups and improve usability in global healthcare systems.

5. Mobile and Web Deployment

Developing a mobile and web-based application will enhance accessibility and enable real-time usage in telemedicine platforms.

6. Large-Scale Evaluation

Future research should include testing on larger and real-world datasets to evaluate scalability, robustness, and performance under heavy workloads.

7. Drug Interaction Prediction

Future work can include advanced models to predict drug–drug interactions and potential side effects more accurately using deep learning techniques.

8.Integration with IoT and Wearable Devices

Connecting the system with wearable health devices can provide real-time health monitoring data, improving recommendation accuracy and preventive care.

SUMMARY

The proposed medicine recommendation system presents a comprehensive and intelligent approach to improving healthcare decision-making by integrating Natural Language Processing (NLP), deep learning, and medical knowledge representation. By utilizing BERT for context-aware sentiment analysis, the system effectively interprets patient reviews and extracts meaningful insights regarding drug effectiveness and side effects. The incorporation of a medical knowledge graph further ensures that recommendations are clinically valid and safe, addressing one of the major limitations of existing systems.

In addition, the inclusion of a personalization module enhances the relevance of recommendations by considering patient-specific attributes such as age, medical history, and allergies. This results in more accurate and user-centric outcomes compared to traditional methods. The modular architecture of the system allows efficient data processing, scalability, and flexibility for future enhancements.

Experimental results demonstrate significant improvements in accuracy, precision, and overall recommendation quality, highlighting the effectiveness of the proposed approach. Despite certain limitations such as computational requirements and dependency on data quality, the system provides a strong foundation for intelligent healthcare applications.

With further improvements such as real-time processing, large-scale deployment, and enhanced security mechanisms, the proposed system has the potential to be integrated into modern healthcare platforms, contributing to more efficient, reliable, and accessible medical services.

Ethical Considerations and Explainability

The proposed system ensures ethical medical recommendations by incorporating knowledge graph validation to prevent unsafe drug suggestions. It avoids harmful drug interactions and contraindications, ensuring patient safety.

Additionally, the system provides explainable outputs by highlighting key factors influencing recommendations, such as patient condition, drug effectiveness, and interaction constraints. This improves transparency and trust in the system.

CONCLUSION

This paper presents an intelligent medicine recommendation system designed to address the limitations of traditional healthcare decision-making approaches. With the increasing availability of unstructured medical data such as patient reviews and clinical notes, the need for automated and accurate analysis has become essential. The proposed system leverages Natural Language Processing (NLP) and deep learning techniques to extract meaningful insights and generate reliable medicine recommendations.

The system integrates a BERT-based sentiment analysis model to capture contextual information from patient reviews, enabling a more accurate understanding of drug effectiveness and adverse effects. Furthermore, the incorporation of a medical knowledge graph ensures that recommendations are clinically valid by verifying drug–disease relationships and identifying contraindications.

Additionally, the incorporation of a personalization module enhances the relevance of recommendations by considering individual patient attributes such as age, medical history, and allergies. This results in more accurate and patient-specific suggestions, improving overall user satisfaction and healthcare outcomes. The modular architecture of the system ensures scalability, flexibility, and ease of integration with future healthcare technologies.

Experimental evaluation demonstrates that the proposed system significantly improves accuracy, precision, and recommendation quality compared to traditional machine learning approaches. It also reduces response time and provides more relevant and safe recommendations

Despite certain limitations such as high computational requirements and dependency on data quality, the proposed system provides a strong foundation for intelligent and automated medicine recommendation. With future enhancements such as real-time deployment, integration with electronic health records, and improved security mechanisms, the system has the potential to be widely adopted in modern healthcare environments.

Overall, the proposed system offers a reliable, efficient, and scalable solution for medicine recommendation, with strong potential for applications in telemedicine and digital healthcare systems. With further enhancements such as real-time processing, large-scale deployment, and improved security, the system can contribute significantly to the advancement of intelligent healthcare technologies.

The experimental findings validate that the proposed system is both effective and reliable, demonstrating strong potential for real-world deployment.

REFERENCES

1. G. V. R. Reddy, M. Y., and T. B. Reddy, “Transformer-Based Healthcare Recommendation Systems,” Computer Research and Development, 2026. Available: https://journalcrd.org/wp-content/uploads/3-CRD3442.pdf
2. J. Jayapradha et al., “Treatment Recommendation Using BERT Personalization (PharmaBERT),” 2024. Available: https://journals.mmupress.com/index.php/jiwe/article/view/1041
3. E. Hassan et al., “Optimizing Disease Prediction Using BERT and Deep Learning,” Scientific Reports, 2024. Available: https://www.nature.com/articles/s41598-024-51615-5
4. X. Zhang et al., “Health Recommender Systems in Healthcare: A Review,” 2023. Available: https://journals.lww.com/inr/fulltext/2023/02000/the_application_of_health_recommender_systems.7.aspx
5. S. Lee et al., “Clinical Decision Transformer for Treatment Recommendation,” 2023. Available: https://arxiv.org/abs/2302.00612
6. M. Nguyen et al., “ALGNet: Graph-Based Medical Recommendation System,” 2023. Available: https://arxiv.org/abs/2312.08377
7. S. Reddy and M. Nair, “AI-Based Healthcare Systems Using NLP,” IEEE Access, 2022. Available: https://ieeexplore.ieee.org/document/9775318
8. S. Wang et al., “Clinical Information Extraction Using BERT,” 2019. Available: https://arxiv.org/abs/1904.03323
9. J. Lee et al., “BioBERT: A Pre-trained Biomedical Language Representation Model,” 2020. Available: https://arxiv.org/abs/1901.08746
10. K. Huang et al., “ClinicalBERT: Modeling Clinical Notes Using BERT,” 2019. Available: https://arxiv.org/abs/1904.05342
11. K. Liang et al., “Medical Knowledge Assisted Models for Healthcare NLP,” 2023. Available: https://arxiv.org/abs/2312.02496

12. F. Rotmensch et al., “Learning a Health Knowledge Graph from EMR Data,” Scientific Reports, 2017. Available: https://www.nature.com/articles/s41598-017-05778-z
13. S. Zhang et al., “Deep Learning-Based Recommender Systems: A Survey,” ACM Computing Surveys, 2019. Available: https://arxiv.org/abs/1707.07435
14. G. Adomavicius and A. Tuzhilin, “Toward the Next Generation of Recommender Systems,” IEEE Transactions on Knowledge and Data Engineering, 2005. Available: https://ieeexplore.ieee.org/document/1423975
15. Y. Koren et al., “Matrix Factorization Techniques for Recommender Systems,” IEEE Computer, 2009. Available: https://ieeexplore.ieee.org/document/5197422
16. H. Wang et al., “Knowledge Graph Convolutional Networks for Recommender Systems,” 2019. Available: https://arxiv.org/abs/1904.12575
17. S. Min et al., “Knowledge Graphs for NLP: A Survey,” 2020. Available: https://arxiv.org/abs/2003.02320
18. Q. Liu et al., “LLM-Based Medication Recommendation Model,” 2024. Available: https://arxiv.org/abs/2402.02803
19. “A Comprehensive Survey of Recommender Systems (2017–2024),” 2025. Available: https://arxiv.org/html/2407.13699v4
20. D. Rajkomar et al., “Machine Learning in Medicine,” New England Journal of Medicine, 2019. Available: https://www.nejm.org/doi/full/10.1056/NEJMra1814259