

CRGANNC-GD: Graph-Diffusion End-to-End Drug Class Recommendation for Stroke Decision Support from Symptom Profiles

Vanessa Innocenta Lonlac Magolac¹, Ribot Fleury Tene Ceskoutsé², Geh Wilson Ejuh³, Jean Pierre Lienou¹

¹Department of Mathematics and Computer Science, University of Dschang, Dschang, Cameroon

²Ecole Nationale Supérieure Polytechnique, University of Yaoundé I, Yaoundé, Cameroon

³Department of Electrical and Electronic Engineering, University of Bamenda, Bamili, Cameroon

Email: magolacvanessa@gmail.com, fleurytene@gmail.com, gehwilsonejeh@yahoo.fr, jean-pierre.lienou@univ-dschang.org

How to cite this paper: Lonlac Magolac, V.I., Tene Ceskoutsé, R.F., Ejuh, G.W. and Lienou, J.P. (2026) CRGANNC-GD: Graph-Diffusion End-to-End Drug Class Recommendation for Stroke Decision Support from Symptom Profiles. *Journal of Intelligent Learning Systems and Applications*, 18, 77-90.

<https://doi.org/10.4236/jilsa.2026.182005>

Received: January 25, 2026

Accepted: March 8, 2026

Published: March 11, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Stroke remains a leading cause of mortality and disability worldwide, requiring timely therapeutic decisions. Existing content-based drug recommendation approaches often rely on static similarity measures and multi-stage pipelines (e.g., clustering followed by instance-based retrieval), which may limit scalability and generalization. We propose CRGANNC-GD, an end-to-end recommendation framework that models patient profiles and medication classes as a heterogeneous graph and performs representation learning through graph diffusion to capture clinically meaningful relationships between symptoms, comorbidities, demographics, and drug classes. Unlike traditional KNN-centric pipelines, CRGANNC-GD learns a task-optimized latent space and outputs calibrated scores for multiple drug classes (antihypertensive, anticoagulant, fibrate). We further incorporate a ranking-aware loss to directly optimize top- k recommendation quality and introduce an explainability module that highlights influential clinical attributes and graph neighborhoods contributing to each recommendation. Experiments on a stroke dataset with 9691 records demonstrate that CRGANNC-GD improves predictive performance and recommendation quality over classical content-based baselines, while maintaining fast inference suitable for clinical decision support and telemedicine settings. Our results suggest that graph-diffusion representation learning provides a scalable and robust alternative for medication class recommendation in stroke management, with a clear pathway to extend the framework toward novel drug discovery via drug-drug and drug-target similarity integration.

Keywords

Stroke, Medication Recommendation, Clinical Decision Support, Heterogeneous Graph Neural Network, Graph Diffusion, Ranking Optimization

1. Introduction

Stroke remains one of the leading causes of mortality and long-term disability worldwide, imposing a substantial burden on healthcare systems and societies. Timely diagnosis and appropriate therapeutic intervention are critical determinants of patient outcomes, as delayed or sub-optimal treatment can lead to irreversible neurological damage or death [1]. In clinical practice, treatment decisions for stroke patients often involve selecting appropriate drug classes, such as anti-hypertensives, anticoagulants, or lipid-modifying agents, based on complex and heterogeneous patient profiles that include demographic factors, comorbidities, vital signs, and neurological symptoms [2]. As the volume and complexity of clinical data continue to grow, there is an increasing need for intelligent decision support systems capable of assisting healthcare professionals in making fast, accurate, and personalized treatment decisions.

In recent years, machine learning has emerged as a promising approach for clinical prediction and decision support, including applications in stroke diagnosis, risk stratification, and treatment recommendation. Content-based recommendation systems and traditional machine learning classifiers have been explored to support medication selection by identifying similarities between patient profiles and historical cases. However, despite their encouraging results, these approaches suffer from several limitations that restrict their effectiveness in real-world clinical settings. In particular, instance-based methods such as K-nearest neighbors often struggle with scalability, sensitivity to high-dimensional data, and limited generalization to unseen patient profiles. Moreover, many existing models rely on fixed similarity measures and sequential processing pipelines, which may fail to capture the complex interdependencies between clinical features, patient subgroups, and therapeutic options [3].

A major challenge in stroke treatment recommendations lies in the heterogeneity and relational nature of clinical data. Patient characteristics, symptoms, comorbidities, and drug classes are not independent entities but are interconnected through intricate and often non-linear relationships. Conventional machine learning models typically operate on flat feature representations and are therefore limited in their ability to exploit such relational structure. As a result, these models may produce incomplete or biased recommendations, particularly in borderline or complex cases [4]. In addition, many deep learning approaches, while powerful, function as black-box systems, raising concerns regarding interpretability and clinical trust. The lack of transparency in model predictions can hinder adoption in

healthcare environments, where understanding the rationale behind a recommendation is essential for validation and accountability.

In this context, we propose CRGANNC-GD, a novel graph-diffusion-based decision support framework for drug class recommendation in stroke management. CRGANNC-GD builds upon our previous CRGANNC [5] model, which combined clustering techniques and instance-based classification to reduce the search space for drug recommendation. While CRGANNC demonstrated strong performance, its reliance on static clustering and K-nearest neighbor retrieval limits its scalability and ability to capture global clinical context. CRGANNC-GD addresses these limitations by modeling patients, clinical features, and drug classes as a heterogeneous graph and leveraging graph neural networks with diffusion mechanisms to learn task-optimized representations in an end-to-end manner.

The proposed framework integrates heterogeneous graph representation learning to capture both local and global clinical relationships, a diffusion process to propagate contextual information across the graph, and a ranking-oriented optimization strategy to directly enhance recommendation quality. In addition, CRGANNC-GD incorporates attention-based mechanisms that provide interpretability by identifying the most influential clinical factors contributing to each recommendation. This combination enables CRGANNC-GD to deliver accurate, scalable, and transparent drug class recommendations tailored to individual stroke patients, making it suitable for deployment in clinical decision support systems and telemedicine applications.

The main contributions of this work can be summarized as follows:

- 1) **Graph-based clinical modeling:** We introduce a heterogeneous graph formulation that explicitly represents relationships between patients, clinical features, and drug classes, enabling the model to exploit the relational structure of stroke data.
- 2) **End-to-end graph-diffusion learning:** We propose a graph neural network framework augmented with diffusion mechanisms to capture global clinical context and overcome the limitations of instance-based similarity methods.
- 3) **Ranking-based drug class recommendation:** We employ a ranking-oriented optimization strategy that improves top- k recommendation performance, aligning the model output with real-world clinical workflows.
- 4) **Interpretability for clinical trust:** We integrate attention-based attribution to provide transparent insights into the clinical factors driving each recommendation, supporting trust and adoption in healthcare settings.

The remainder of this paper is organized as follows. Section II reviews related work on stroke prediction and treatment recommendation. Section III describes the proposed CRGANNC-GD methodology in detail. Section IV presents the experimental setup, results, and discussion. Finally, Section V concludes the paper and outlines directions for future research.

2. Related Work

Machine learning for stroke decision support has expanded rapidly in recent

years, covering a wide range of applications including risk prediction, prognosis estimation, and clinical decision assistance. In both acute management and secondary prevention settings, stroke care is inherently time-sensitive and relies on heterogeneous clinical data such as demographic information, comorbidities, vital signs, neurological symptoms, laboratory markers, and medication histories. Several recent surveys and position papers highlight the significant potential of artificial intelligence to enhance stroke care, while also emphasizing critical challenges related to clinical validation, safety, interpretability, and real-world deployment in hospital environments [6]-[8]. Despite these advances, the majority of existing studies primarily focus on stroke risk assessment or outcome prediction, whereas comparatively fewer works directly address medication or therapy recommendation, even though such decisions play a central role in clinical workflows.

A fundamental limitation of conventional clinical prediction pipelines lies in their reliance on flat feature representations, where each patient is encoded as a fixed-length vector and independent mappings are learned between features and outcomes. This paradigm often fails to capture the rich relational structure inherent in healthcare data, where patients, symptoms, diagnoses, laboratory tests, and drug classes are interconnected through complex and non-linear relationships. Recent methodological reviews increasingly advocate for graph-based modeling as a natural and powerful paradigm for electronic health records, as it enables the explicit representation of heterogeneous entities and their interactions, leading to improved representation learning and more clinically meaningful reasoning [9]-[11].

Within this broader context, medication recommendation has traditionally been approached using sequential deep learning models, such as recurrent neural networks or Transformer-based architectures, trained on longitudinal patient admissions. These approaches are often augmented with drug-drug interaction constraints to mitigate safety risks. A comprehensive review of deep learning techniques for medication recommendation highlights persistent challenges including data sparsity, temporal dependency modeling, and the integration of drug interaction knowledge [12]. To address these limitations, recent studies have begun to incorporate knowledge graphs and graph neural networks into medication recommendation frameworks. For example, Mishra and Shridevi proposed a knowledge-graph-driven framework that constructs admission-level clinical and medication graphs, combines longitudinal modeling with graph neural networks, and integrates external ontologies and drug interaction knowledge to improve recommendation safety and accuracy [13]. Similarly, Liu *et al.* introduced a multi-granularity drug recommendation model that incorporates drug-drug interaction-aware regularization to balance predictive performance and clinical safety [14]. Other recent approaches further explore graph encoders and attention mechanisms to enhance medication recommendation in sparse and heterogeneous EHR settings, demonstrating consistent performance gains over purely sequential or feature-based models [15] [16].

Beyond medication recommendation, graph neural networks have been increasingly adopted for a wide range of clinical prediction tasks, including patient risk stratification and disease progression modeling. Frameworks such as GraphCare leverage external biomedical knowledge graphs and patient-specific subgraphs to enhance downstream healthcare prediction tasks, including drug recommendation, across multiple EHR benchmarks [16]. Recent reviews in clinical informatics provide systematic overviews of GNN-based approaches for healthcare applications, discussing common graph construction strategies (e.g., patient graphs, visit graphs, and medical code graphs) as well as open challenges related to robustness, calibration, and explainability [9]. In the stroke domain, emerging studies have explored graph-based models for clinically relevant prediction problems, such as atrial fibrillation risk in critically ill stroke patients, and emphasize the importance of explainability tools—including attention mechanisms and attribution analysis—to foster clinical trust and adoption [17].

Despite these methodological advances, stroke-specific medication recommendations remain relatively underexplored compared to the broader literature on general medication recommendations. Our previous work introduced CRGANNC, a content-based recommendation framework that combines Gaussian Mixture Models, Affinity Propagation, and K-nearest neighbors to reduce the search space and recommend stroke drug classes, including antihypertensive, anticoagulant, and fibrate therapies, based on patient symptom profiles [5]. While CRGANNC demonstrated strong predictive precision and improved computational efficiency, its reliance on static clustering and instance-based retrieval limits scalability and constrains its ability to capture global clinical context, particularly when interactions between symptoms, comorbidities, and drug classes are complex or non-local.

Taken together, the existing literature reveals two converging trends: a shift from flat, instance-based learning toward graph-structured representations that explicitly encode clinical relationships, and an increasing emphasis on safe, interpretable recommendation systems guided by medical knowledge and interaction constraints [13] [14] [16]. However, there remains a clear gap for stroke-oriented drug class recommendation frameworks that simultaneously offer scalability beyond KNN-style retrieval, the ability to model global clinical context through diffusion or long-range information propagation, and interpretability suitable for real-world clinical adoption. The proposed CRGANNC-GD framework is designed to address this gap by learning end-to-end heterogeneous graph representations of patients, clinical features, and drug classes, integrating diffusion-based context propagation and ranking-oriented optimization to align with practical stroke decision-making workflows.

3. Methods

This work presents CRGANNC-GD (**Figure 1**), an end-to-end graph-based decision support model for drug class recommendation in stroke management.

CRGANNC-GD extends the original CRGANNC framework by replacing its multi-stage clustering and instance-based retrieval pipeline with heterogeneous graph representation learning and graph diffusion, enabling scalable, robust, and clinically interpretable recommendations. The proposed methodology is organized around three main contributions, detailed below.




Figure 1. Overview of the proposed CRGANNC-GD framework for drug class recommendation in stroke management. The model integrates clinical graph construction with SMOTE-based augmentation, heterogeneous graph representation learning with diffusion, and ranking-based drug class recommendation with interpretability.

First contribution. The first contribution is a clinical graph construction and data augmentation prototype designed to address data imbalance and sparsity commonly observed in stroke datasets. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denote a dataset of stroke patients, where $x_i \in \mathbb{R}^d$ represents clinical features (symptoms, demographics, biomarkers) and $y_i \in \{0, 1\}^C$ denotes prescribed drug classes, with $C = 3$ (antihypertensive, anticoagulant, fibrate). From this dataset, a heterogeneous clinical graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is constructed, comprising patient nodes P , feature nodes F , and drug class nodes D . Edges are defined between patient and feature nodes when a feature is clinically relevant, and between patient and drug class nodes when a prescription is observed. Edge weights are computed as:

$$w_{ij} = \phi(x_{ij}), \quad (1)$$

where $\phi(\cdot)$ denotes normalized feature intensity.

To mitigate class imbalance among drug classes, a SMOTE-based oversampling strategy is applied prior to graph construction. Given a minority patient sample x_i and one of its k nearest neighbors x_j , a synthetic patient sample is generated as:

$$x_{new} = x_i + \lambda(x_j - x_i), \quad \lambda \in [0, 1]. \quad (2)$$

The resulting synthetic samples are incorporated as new patient nodes, and corresponding edges are added while preserving the original graph topology.

Algorithm 1. Clinical graph augmentation.

- 1: Load patient dataset $\mathcal{D}$
 - 2: Handle missing values using imputation
 - 3: Apply SMOTE oversampling on minority drug classes
 - 4: Construct heterogeneous graph $\mathcal{G} = (P, F, D, E)$
 - 5: Assign edge weights using Equation (1)
 - 6: Save augmented graph $\mathcal{G}_{aug}$
-

Second contribution. The second contribution of CRGANNC-GD is a heterogeneous graph representation learning model with diffusion, designed to capture both local and global clinical context. Given the augmented graph $\mathcal{G}_{aug}$, each node $v \in \mathcal{V}$ is initialized as:

$$h_{p_i}^{(0)} = W_p x_i, \quad h_{f_j}^{(0)} = e_{f_j}, \quad h_{d_c}^{(0)} = e_{d_c}, \quad (3)$$

where W_p is a learnable projection matrix and $e_{(\cdot)}$ are trainable embeddings. Representation learning is performed using a relation-aware message passing mechanism inspired by heterogeneous graph neural networks. At layer l , messages are computed as:

$$m_{v,r}^{(l)} = \sum_{u \in \mathcal{N}_r(v)} \alpha_r^{(l)}(u, v) W_r^{(l)} h_u^{(l)}, \quad (4)$$

and node embeddings are updated as:

$$h_v^{(l+1)} = \sigma \left(W_0^{(l)} h_v^{(l)} + \sum_{r \in \mathcal{R}} m_{v,r}^{(l)} \right), \quad (5)$$

where $\alpha_r^{(l)}(u, v)$ denotes attention coefficients and $\sigma(\cdot)$ is a nonlinear activation function.

To further incorporate global structural information, a diffusion process is applied to the learned embeddings:

$$Z = (1 - \beta) \sum_{k=0}^K \beta^k \tilde{A}^k H, \quad (6)$$

where $\tilde{A}$ is the normalized adjacency matrix, H denotes the stacked node embeddings, and β controls the diffusion depth.

Algorithm 2. Heterogeneous graph representation learning.

- 1: Initialize node embeddings using Equation (3)
 - 2: **for** $l = 1$ to L **do**
 - 3: Aggregate relation-aware messages using Equation (4)
 - 4: Update node embeddings using Equation (5)
 - 5: **end for**
 - 6: Apply graph diffusion using Equation (6)
 - 7: Extract final embeddings Z
-

Third contribution. The third contribution of CRGANNC-GD is a ranking-based drug class recommendation mechanism built upon the learned graph embeddings. For each patient-drug class pair (p_i, d_c) , a relevance score is computed as:

$$\hat{y}_{i,c} = \sigma(\mathbf{z}_{p_i}^\top \mathbf{z}_{d_c}), \quad (7)$$

where $\mathbf{z}_{(\cdot)}$ denotes diffused embeddings. Drug classes are ranked according to these scores, enabling top- k recommendations.

Model training is performed using a joint classification and ranking objective:

$$\mathcal{L} = \lambda_1 \text{BCE}(y, \hat{y}) + \lambda_2 \sum_{(c^+, c^-)} \log(1 + \exp(-(s_{c^+} - s_{c^-}))), \quad (8)$$

which encourages accurate prediction while directly optimizing recommendation quality.

Algorithm 3. CRGANNC-GD drug recommendation.

-
- 1: Input patient profile x_i
 - 2: Project x_i into the graph embedding space
 - 3: Compute relevance scores for all drug classes
 - 4: Rank drug classes by score
 - 5: Return Top- k recommended drug classes
-

Finally, to enhance clinical interpretability, CRGANNC-GD quantifies feature importance using attention-based attribution:

$$\text{Imp}(f_j \rightarrow p_i) = \sum_{l,r} \alpha_r^{(l)}(f_j, p_i) \|\mathbf{W}_r^{(l)} \mathbf{h}_{f_j}^{(l)}\|. \quad (9)$$

This enables healthcare professionals to identify the most influential symptoms and clinical factors underlying each recommendation. By integrating clinical graph construction, graph-diffusion representation learning, and ranking-based recommendation, CRGANNC-GD overcomes the scalability and generalization limitations of traditional content-based systems and provides an effective decision support framework for stroke care and telemedicine applications.

4. Experimental Results and Discussion

To train, evaluate, and validate the proposed CRGANNC-GD model, we used a stroke dataset composed of 9691 patient records, including 1206 stroke cases and 8485 healthy individuals. This dataset was constructed following the same preprocessing pipeline used in our previous work on CRGANNC in order to ensure comparability and reproducibility. Clinical attributes included demographic information, comorbidities, vital signs, and stroke-related symptoms such as weakness, dizziness, headache, nausea, and vomiting. Drug classes considered in this study were antihypertensive, anticoagulant, and fibrate medications. The dataset was split into training (70%), validation (15%), and test (15%) sets using stratified sampling to preserve the original class distribution. To avoid data leakage, the da-

taset is first split into training, validation, and test sets using stratified sampling. SMOTE oversampling is then applied exclusively to the training set, while validation and test sets remain untouched. Continuous features are represented as shared feature nodes with weighted edges proportional to normalized feature values. Missing values are handled using SimpleImputer with median strategy (for continuous variables) prior to graph construction. By relying on the same publicly available data source and preparation strategy, we ensure a fair comparison between CRGANNC-GD and existing content-based and machine learning baselines.




Figure 2. Recall comparison between CRGANNC-GD and baseline models.

To isolate the contribution of the diffusion and ranking components, we additionally include a standard GAT/R-GCN baseline using the same graph structure but without diffusion or ranking-aware loss. Healthy individuals are included as graph nodes to enrich representation learning and facilitate message passing across shared clinical features. They are excluded from the ranking loss computation, which is applied only to stroke patients. We compared the performance of CRGANNC-GD against several competitive models, including the original CRGANNC framework, LightGBM, Multilayer Perceptron (MLP), and TabTransformer (**Figure 2**). CRGANNC relies on Gaussian Mixture Models, Affinity Propagation, and K-nearest neighbors to reduce the search space and recommend drug classes, while the other baselines represent strong non-graph machine learning approaches com-

monly used in clinical prediction tasks. Performance was evaluated using macro-averaged Precision, Recall, and F1-score, as well as recommendation-oriented metrics such as Hit@1, Hit@3, and NDCG@3. All experiments were repeated across five independent runs, and average results are reported.

The quantitative results demonstrate that CRGANNC-GD consistently outperforms all baseline models across both classification and recommendation metrics. In particular, CRGANNC-GD achieved a macro F1-score of 0.98, compared to 0.96 for CRGANNC and 0.95 for LightGBM and TabTransformer (**Table 1**). This improvement highlights the benefit of learning patient representations through heterogeneous graph modeling rather than relying on fixed similarity measures. The superior Hit@1 and Hit@3 scores obtained by CRGANNC-GD further confirm its effectiveness in ranking clinically relevant drug classes, which is essential for real-world decision support scenarios where multiple therapeutic options must be considered.

Table 1. Overall performance comparison between CRGANNC-GD and baseline models.

Model	Precision	Recall	F1-score	Hit@1	Hit@3
CRGANNC	0.98	0.95	0.96	0.90	0.97
LightGBM	0.96	0.94	0.95	0.89	0.96
MLP	0.95	0.93	0.94	0.87	0.95
TabTransformer	0.96	0.94	0.95	0.88	0.96
CRGANNC-GD	0.99	0.97	0.98	0.92	0.99

To assess the robustness and stability of CRGANNC-GD during training, we monitored the training loss over multiple epochs and across five independent runs. In each run, the model was trained for 10 epochs, and the mean loss was recorded after each epoch (**Figure 3**). The observed loss curves exhibit a consistent decreasing trend, indicating effective convergence of the model. Early epochs show a rapid reduction in loss, followed by a gradual stabilization, suggesting that the model successfully learns meaningful latent representations of patients, features, and drug classes. The limited variance observed between different runs confirms the reliability and repeatability of the training process. These results indicate that CRGANNC-GD does not suffer from instability or overfitting despite the increased model complexity introduced by graph representation learning and diffusion.

We further conducted ablation experiments to evaluate the contribution of key components of the proposed framework (**Table 2**). Removing the diffusion mechanism resulted in a noticeable degradation of performance, particularly in recall and Hit@3 metrics, indicating that global context propagation plays a critical role in capturing long-range clinical dependencies. Similarly, replacing the ranking-based loss with a pure classification objective led to reduced recommendation quality, demonstrating the importance of directly optimizing for ranking in drug

recommendation tasks. These findings confirm that each component of CRGANNC-GD contributes independently and synergistically to the observed performance gains.




Figure 3. Training loss evolution of CRGANNC-GD over 10 epochs across five independent runs.

Table 2. Ablation study evaluating the contribution of CRGANNC-GD components.

Model Variant	Precision	Recall	F1-score	Hit@3
CRGANNC-GD (full)	0.99	0.97	0.98	0.99
w/o diffusion	0.97	0.94	0.95	0.96
w/o ranking loss	0.97	0.93	0.95	0.95
w/o attention	0.96	0.93	0.94	0.95

From a clinical perspective, the improved recall achieved by CRGANNC-GD is particularly significant, as it reduces the likelihood of missing clinically relevant drug classes for stroke patients. In time-sensitive conditions such as stroke, where delayed or inappropriate treatment can lead to severe outcomes, such improvements can have a direct impact on patient care. Moreover, the ranking-based recommendation output aligns well with real-world clinical workflows, where healthcare professionals typically evaluate several treatment options rather than relying on a

single deterministic prediction.

An additional strength of CRGANNC-GD lies in its interpretability. By leveraging attention-based attribution within the heterogeneous graph, the model provides insights into the most influential symptoms and clinical features driving each recommendation. This transparency enhances trust and usability, which are critical factors for the adoption of artificial intelligence systems in clinical settings. Unlike traditional black-box classifiers, CRGANNC-GD enables healthcare professionals to better understand and contextualize the model's recommendations.

Overall, the experimental results demonstrate that CRGANNC-GD represents a significant advancement over traditional content-based recommendation systems such as CRGANNC, as well as over strong machine learning baselines. By integrating heterogeneous graph representation learning, diffusion-based global context modeling, and ranking-oriented optimization, the proposed framework provides a robust, scalable, and interpretable solution for drug class recommendation in stroke management. These findings highlight the potential of graph-based decision support systems to enhance clinical decision-making and support telemedicine applications, particularly in resource-constrained or underserved settings.

5. Conclusion and Future Work

In this study, we proposed CRGANNC-GD, an end-to-end graph-diffusion framework for drug class recommendation in stroke management. Unlike traditional content-based recommendation systems that rely on fixed similarity measures and multi-stage pipelines, CRGANNC-GD models patients, clinical features, and drug classes within a heterogeneous graph and learns task-specific representations through relation-aware message passing. By integrating a diffusion mechanism, the framework captures global clinical context beyond immediate neighborhoods, while a ranking-oriented objective directly optimizes top- k recommendation quality. Experimental results show that CRGANNC-GD consistently improves both classification metrics and recommendation performance compared to the original CRGANNC model and strong machine learning baselines. In addition, attention-based attribution provides clinically meaningful explanations by highlighting influential symptoms and factors driving each recommendation.

Future work will focus on several directions. First, we plan to extend the model to recommend drugs that were not initially administered to a patient by integrating drug-drug similarity, pharmacological knowledge, and drug-target relationships into the graph. Second, we will investigate calibration and uncertainty estimation to support safer clinical deployment, especially in borderline cases where multiple therapeutic options may be appropriate. Third, we aim to evaluate CRGANNC-GD on external cohorts and real-world clinical datasets to validate generalization across populations and healthcare settings. Finally, incorporating causal inference and counterfactual reasoning could enable the model to move

beyond association-based recommendations toward clinically actionable insights for precision medicine and treatment optimization in stroke care.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Conflicts of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] de Andrade, J.B.C., Fagundes, T.P., Katsuyama, E. and Silva, G.S. (2026) Digital Health in Low-Resource Settings: Comprehensive Challenges and Opportunities with a Focus on Stroke Care. *Stroke*, **57**, 245-253.
<https://doi.org/10.1161/strokeaha.125.050448>
- [2] Nayyar, S. and Brown, P. (2026) Medications for Anti-Coagulation in Stroke Prevention. In: Thomas, R.E., Ed., *Medications for Older Persons*, Elsevier, 141-154.
<https://doi.org/10.1016/b978-0-443-29025-1.00005-9>
- [3] Kücking, F., Busch, D.A., Przysucha, M., Kutza, J., Hannemann, N., Hüsters, J., *et al.* (2026) Impact of AI Recommendation Correctness on Diagnostic Accuracy in Clinical Decision-Making. *International Journal of Medical Informatics*, **207**, Article ID: 106223. <https://doi.org/10.1016/j.ijmedinf.2025.106223>
- [4] Ayata, C., Bath, P.M., Planas, A.M., Allan, S.M., Boltze, J., Cabeen, R.P., *et al.* (2026) Preclinical Ischemic Stroke Multicenter Trials (PRISM) Collective Statement: Opportunities, Challenges, and Recommendations for a New Era. *Stroke*, **57**, e12-e40.
<https://doi.org/10.1161/strokeaha.125.052056>
- [5] Tene Ceskoutse, R.F., *et al.* (2024) Crgann: A Clustering-Based Content Recommendation Model for Drug Class Detection in Stroke Patients. *Expert Systems with Applications*.
- [6] Liu, Y., Wen, Z.J., Wang, Y.R., Zhong, Y.X., Wang, J.X., Hu, Y.H., Zhou, P. and Guo, S.M. (2024) Artificial Intelligence in Ischemic Stroke Images: Current Applications and Future Directions. *Frontiers in Neurology*, **15**, Article 1418060.
<https://doi.org/10.3389/fneur.2024.1418060>
- [7] Amann, J. (2021) Machine Learning (ML) in Stroke Medicine: Opportunities and Challenges for Risk Prediction and Prevention. In: Jotterand, F. and Ienca M., Eds., *Artificial Intelligence in Brain and Mental Health: Philosophical, Ethical & Policy Issues*, Springer International Publishing, 57-71.
https://doi.org/10.1007/978-3-030-74188-4_5
- [8] Basem, J., Mani, R., Sun, S., Gilotra, K., Dianati-Maleki, N. and Dashti, R. (2025) Clinical Applications of Artificial Intelligence and Machine Learning in Neurocardiology: A Comprehensive Review. *Frontiers in Cardiovascular Medicine*, **12**, Article 1525966. <https://doi.org/10.3389/fcvm.2025.1525966>
- [9] Hancox, Z., Pang, A., Conaghan, P.G., Kingsbury, S.R., Clegg, A. and Relton, S.D. (2024) A Systematic Review of Networks for Prognostic Prediction of Health Outcomes and Diagnostic Prediction of Health Conditions within Electronic Health Rec-

- ords. *Artificial Intelligence in Medicine*, **158**, Article 102999. <https://doi.org/10.1016/j.artmed.2024.102999>
- [10] Vaida, M. and Huang, Z.Y. (2026) Multimodal Graph Neural Networks in Healthcare: A Review of Fusion Strategies across Biomedical Domains. *Frontiers in Artificial Intelligence*, **8**, Article 1716706. <https://doi.org/10.3389/frai.2025.1716706>
 - [11] Al Khatib, H.S., Neupane, S., Manchukonda, H.K., Golilarz, N.A., Mittal, S., Amirlatif, A. and Rahimi, S. (2024) Patient-Centric Knowledge Graphs: A Survey of Current Methods, Challenges, and Applications. *Frontiers in Artificial Intelligence*, **7**, Article 1388479. <https://doi.org/10.3389/frai.2024.1388479>
 - [12] Ali, Z., Huang, Y., Ullah, I., Feng, J.L., Deng, C., Thierry, N., *et al.* (2022) Deep Learning for Medication Recommendation: A Systematic Survey. *Data Intelligence*, **5**, 303-354. https://doi.org/10.1162/dint_a_00197
 - [13] Mishra, P. and Shridevi, S. (2024) Knowledge Graph-Driven Medicine Recommendation System Using Graph Neural Networks on Longitudinal Medical Records. *Scientific Reports*, **14**, Article 22611. <https://doi.org/10.1038/s41598-024-75784-5>
 - [14] Liu, F., Wang, W., Zheng, J., Xie, Y., Wang, X. and Zhang, D. (2025) EDRMM: Enhancing Drug Recommendation via Multi-Granularity and Multi-Attribute Representation. *BMC Bioinformatics*, **26**, Article No. 173. <https://doi.org/10.1186/s12859-025-06167-4>
 - [15] Wang, Y.G. and Lu, B.X. (2024) Drug Recommendation Combining Attention Mechanism and Graph Encoder. *Proceedings of the 2023 International Conference on Information Education and Artificial Intelligence*, New York, 22-24 December 2023, 750-755. <https://doi.org/10.1145/3660043.3660177>
 - [16] Jiang, P.C., Xiao, C., Cross, A. and Sun, J.M. (2024) Graphcare: Enhancing Healthcare Predictions with Personalized Knowledge Graphs.
 - [17] Rivera-Juzga, C.Y., Mistry, D., Knowles, A.T., Lip, G.Y.H., Ortega-Martorell, S. and Olier, I. (2025) Graph Neural Networks for the Enhanced Prediction of New Atrial Fibrillation Episodes after Stroke. *Computers in Biology and Medicine*, **197**, Article 111095. <https://doi.org/10.1016/j.compbiomed.2025.111095>