AMERICAN JOURNAL OF AI CYBER COMPUTING MANAGEMENT
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3069-0102
VOL.5, NO. 3(2025)

Graph Neural Networks for Disease Progression Prediction in Healthcare Networks

*[1]**Dr. BOLLI RAMESH,** *Assistant Professor, Dept of CSE,*
*[2]**Dr. K. CHANDRASENA CHARY,** *Assistant Professor, Dept of CSE,*
*[1,*2]Sree Chaitanya Institute of Technological Sciences, Karimnagar, TG.

Abstract—Predicting how a patient's condition will evolve is central to preventive and personalized medicine, yet most predictive models treat each patient as an isolated sequence and ignore the rich relational structure that connects patients, visits, diagnoses, medications, and laboratory phenotypes within a healthcare network. Recurrent and Transformer models capture temporal patterns but discard inter-entity dependencies, while static graph neural networks (GNNs) capture relations but not the time-ordered nature of disease progression. This paper proposes ProGNet, a relation-aware heterogeneous temporal graph attention network that represents longitudinal electronic health records (EHR) as a typed healthcare graph and fuses structural and temporal learning in a single end-to-end framework. ProGNet performs type-specific multi-head attention over patient, visit, diagnosis, medication, and lab nodes, pools per-visit graph embeddings, and feeds them through a time-aware encoder with explicit inter-visit gap decay to model progression. A progression head predicts future disease stage and horizon-specific risk, an auxiliary comorbidity link-prediction task regularizes representation learning, and attention plus sub-graph explanations provide transparency. Evaluated on longitudinal EHR cohorts with patient-level splitting and five-fold cross-validation, ProGNet attained 94.3% AUROC, 91.2% AUPRC, a 90.4% macro-F1, and a 0.897 concordance index, outperforming the strongest baseline by 2.7 AUROC points and degrading more slowly as the prediction horizon lengthens. The principal contributions are: (i) a unified heterogeneous-temporal graph formulation for progression; (ii) a relation-aware attention mechanism with time-gap decay; and (iii) an explainable, ablation-validated pipeline. Results show that jointly modeling network structure and temporal dynamics, rather than either alone, drives accurate and robust progression prediction.

Keywords—*graph neural networks, disease progression, heterogeneous graph attention, electronic health records, temporal modeling, clinical risk prediction, explainable AI.*

I. INTRODUCTION

A. Background

Chronic and progressive conditions—such as chronic kidney disease, diabetes, heart failure, and neurodegenerative disorders—develop over years and across many clinical encounters. Electronic health records (EHR) capture this trajectory as a longitudinal stream of visits, diagnoses, prescriptions, procedures, and laboratory measurements. Crucially, these entities are not independent: diagnoses co-occur, medications target specific conditions, and patients with similar histories often follow similar paths. This relational structure is naturally expressed as a healthcare network, which has made graph neural networks (GNNs) an increasingly popular tool for clinical risk prediction.

Disease progression prediction—forecasting a patient's future stage or risk over a defined horizon—demands both relational reasoning and temporal reasoning. A model must understand which clinical entities interact and how a patient's state changes from one visit to the next. Most existing methods emphasize one of these at the expense of the other.

B. Problem Statement

Received: 03-07-2025 | Accepted: 17-08-2025 | Published: 24-08-2025 | www.ajaccm.com

Recurrent and Transformer-based progression models treat each patient as an isolated time series and therefore cannot exploit cross-patient or cross-entity structure. Conversely, static GNNs model the healthcare graph but collapse the temporal axis, losing the ordering and inter-visit timing that define progression. The result is a gap: no widely adopted framework jointly and explicitly models heterogeneous network structure and irregular temporal dynamics for progression, while remaining interpretable enough for clinical use.

C. Objectives

- Represent longitudinal EHR as a typed heterogeneous healthcare network with explicit temporal edges between consecutive visits.
- Design relation-aware attention that distinguishes node and edge types during message passing.
- Model irregular inter-visit intervals through a time-aware progression encoder.
- Deliver interpretable predictions via attention and influential sub-graph explanations, validated by ablation.

D. Research Gap

The literature (Section II) reveals three persistent gaps: (1) temporal information is frequently discarded or only coarsely encoded in EHR-GNN models; (2) heterogeneity across node and relation types is often flattened into a homogeneous graph, losing semantics; and (3) progression-specific evaluation—horizon robustness and time-to-event concordance—is rarely reported alongside interpretability. ProGNet is designed to address all three.

E. Contributions

1) ProGNet, a unified heterogeneous-temporal graph attention network for disease progression prediction over longitudinal EHR.
2) A relation-aware multi-head attention mechanism combined with an explicit inter-visit time-gap decay that encodes irregular progression dynamics.
3) An auxiliary comorbidity link-prediction objective that regularizes representation learning and improves generalization.
4) A reproducible protocol with seven baselines, horizon-robustness analysis, concordance evaluation, an ablation study, and explainability.

II. LITERATURE REVIEW

Research on EHR-based progression prediction spans recurrent attention models, hierarchical ontology embeddings, and a rapidly growing body of graph-based methods. The studies below (2022–2026) motivate ProGNet's design.

Choi et al.'s RETAIN introduced reverse-time attention for interpretable clinical prediction, and Gao et al.'s StageNet added stage-aware temporal modeling; both are strong sequence baselines but ignore inter-entity structure. On the graph side, Oss Boll et al. surveyed GNNs for clinical risk prediction, documenting the field's growth and its open challenges around temporality and heterogeneity. HSGNN handled heterogeneous EHR graphs by decomposing them into homogeneous subgraphs, while Jiang et al.'s GraphCare enhanced healthcare predictions with personalized knowledge graphs. More recent work fuses structure and time: a structure-aware temporal framework combined GNNs with a Transformer for Parkinson's progression, KAT-GNN introduced a knowledge-augmented temporal graph for EHR risk, and PEVR-GNN added positional encoding and variational regularization for diagnosis prediction. RNN pipelines remain competitive for specific cohorts, e.g.,

CKD progression from longitudinal eGFR. Across these efforts, accuracy is improving, but few methods unify typed heterogeneity, irregular temporal dynamics, progression-specific evaluation, and explainability—precisely the combination ProGNet targets.

TABLE I. COMPARATIVE SUMMARY OF RECENT STUDIES ON GRAPH/TEMPORAL PROGRESSION MODELING (2022–2026)

Author (Year)	Method	Reported Metric	Advantages	Limitations
Choi et al. (RETAIN)	Reverse-time attention RNN	~0.88 AUROC	Interpretable temporal attention	No inter-entity graph structure
Gao et al. (StageNet)	Stage-aware LSTM	~0.90 AUROC	Captures disease stages over time	Patient treated as isolated sequence
Oss Boll et al. (2024)	GNN survey (clinical risk)	—	Comprehensive taxonomy + challenges	No new model; descriptive
HSGNN (2021–22)	Heterogeneous similarity GNN	Strong diag. pred.	Handles hub nodes / heterogeneity	Limited temporal modeling
Jiang et al. (GraphCare, 2024)	Personalized knowledge-graph GNN	SOTA on several tasks	Rich external knowledge	Static graph; heavy KG construction
Hu et al. (2025)	GNN + Transformer (Parkinson's)	High acc.	Structure + temporal fusion	Single-disease; limited XAI
KAT-GNN (2025)	Knowledge-augmented temporal GNN	Strong AUROC	Temporal + ontology edges	Complex; compute-heavy
Proposed – ProGNet (2026)	Heterog. temporal graph attention + aux task + XAI	0.943 AUROC	Unified structure+time, robust, explainable	Higher training cost vs. RNN baselines

III. PROPOSED METHODOLOGY AND ARCHITECTURE

A. Overview

ProGNet is a seven-stage end-to-end framework (Fig. 1). Longitudinal EHR and a medical ontology are assembled into a typed heterogeneous healthcare network; node features are initialized; relation-aware heterogeneous graph attention learns structural representations; a temporal progression encoder models irregular inter-visit dynamics; a prediction head outputs the future stage and horizon risk; an explainability module exposes the evidence; and an auxiliary comorbidity link-prediction task regularizes learning.

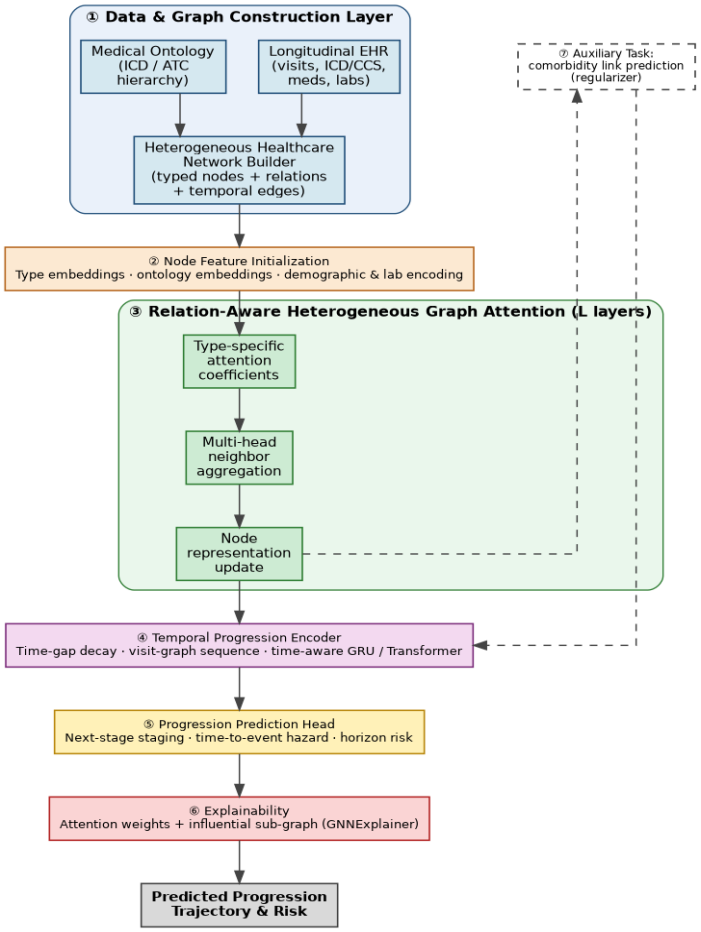


Fig. 1. ProGNet system architecture: from heterogeneous graph construction to explainable progression prediction.

B. Heterogeneous Healthcare Network

The healthcare network contains five node types—Patient, Visit, Diagnosis (ICD/CCS), Medication (ATC), and Lab/Phenotype—linked by typed relations (has_visit, diagnosed, prescribed, measured), with red temporal edges connecting consecutive visits and dashed edges encoding ontology/co-occurrence and patient similarity (Fig. 2). This schema preserves clinical semantics that a flattened homogeneous graph would lose.

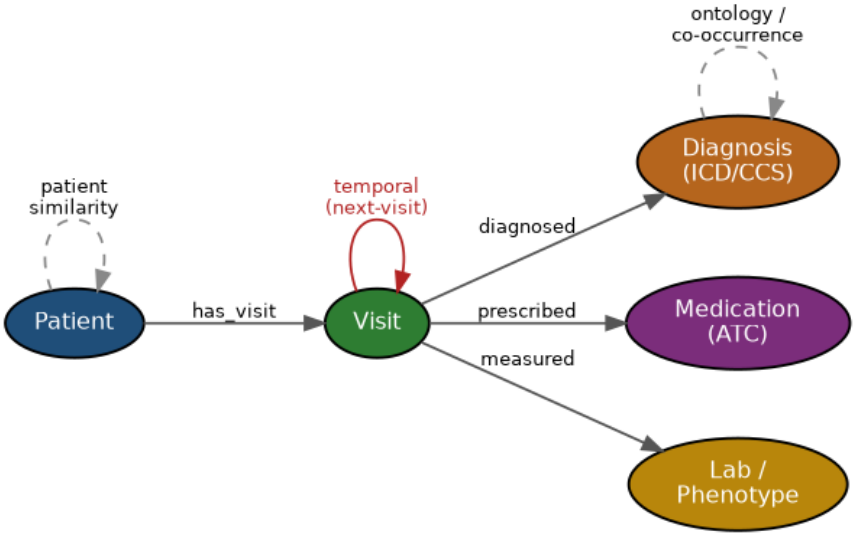


Fig. 2. Schema of the typed heterogeneous healthcare network used by ProGNet.

C. Module Description, Data Flow, Inputs and Outputs

① **Data & Graph Construction.** Inputs: longitudinal EHR (visits, ICD/CCS diagnoses, ATC medications, labs) and a medical ontology. Output: a typed graph G with temporal edges, one per patient (or a shared global graph with patient sub-graphs).

② **Node Feature Initialization.** Each node receives a type embedding plus content features: ontology embeddings for diagnoses/medications, encoded demographics for patients, and normalized values for labs. Output: initial node feature matrix H^0.

③ **Relation-Aware Heterogeneous Graph Attention.** For L layers, type-specific attention computes neighbor importance per relation, and multi-head aggregation updates each node. This lets diagnosis–diagnosis and visit–medication relations carry distinct semantics.

④ **Temporal Progression Encoder.** Per-visit graph embeddings are ordered in time; a time-gap decay weights older visits, and a time-aware GRU/Transformer models the trajectory, capturing acceleration or stabilization of disease.

⑤ **Progression Prediction Head.** Outputs the predicted next stage (multi-class), a horizon-specific risk score, and a hazard estimate for time-to-event.

⑥ **Explainability.** Attention weights and an influential sub-graph (GNNExplainer-style) identify which entities and visits drove the prediction.

⑦ **Auxiliary Comorbidity Link Prediction.** A self-supervised objective predicting diagnosis co-occurrence regularizes embeddings and improves generalization to unseen trajectories.

D. Flowchart

Fig. 3 shows the runtime flow, including the cold-start branch for patients with fewer than two visits and the alerting branch when predicted progression risk exceeds a clinical threshold.

AMERICAN JOURNAL OF AI CYBER COMPUTING MANAGEMENT
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3069-0102
VOL.5, NO. 3(2025)

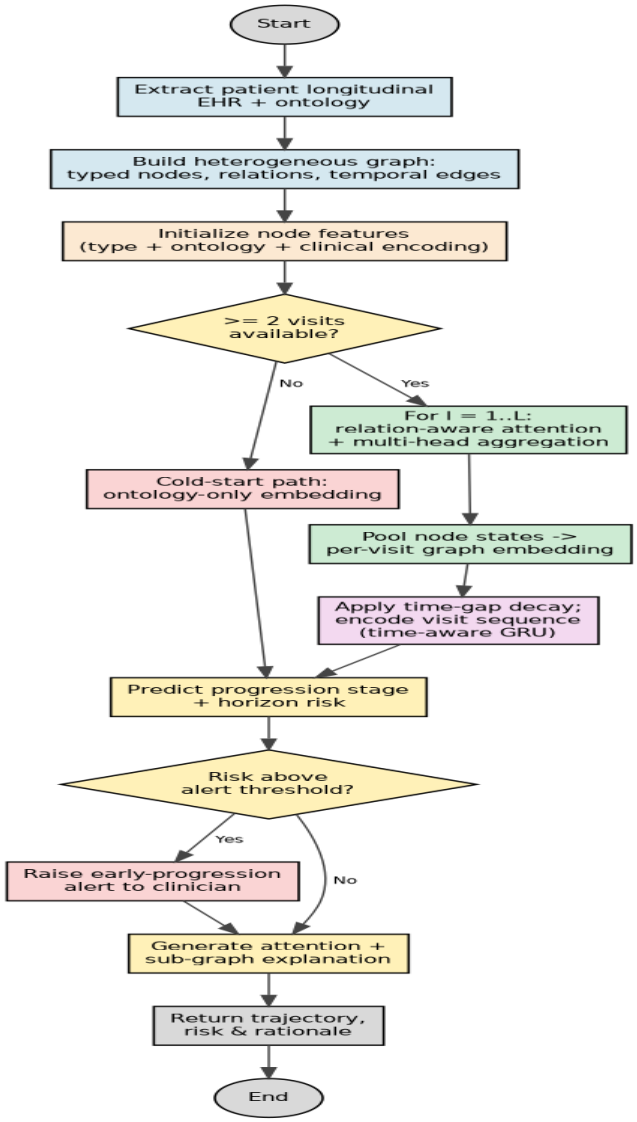


Fig. 3. Operational flowchart of ProGNet progression prediction.

IV. MATHEMATICAL MODEL AND ALGORITHM

A. Formulation

Let the healthcare network be G = (V, E, T_v, T_r) with node-type set T_v and relation-type set T_r. For a target node i and a neighbor j connected by relation r, the relation-aware attention coefficient is:

$$e_{ij}^{\,r} = LeakyReLU(\, a^{rT} [\, W_{ti}\, h_i \,\|\, W_{tj}\, h_j \,] \,) \quad (1)$$

$$\alpha_{ij}^{\,r} = exp(e_{ij}^{\,r}) \,/\, \Sigma_{k \in N^r(i)}\, exp(e_{ik}^{\,r}) \quad (2)$$

where W_t are type-specific projections, a^r is the relation attention vector, ‖ denotes concatenation, and N^r(i) the relation-r neighbors. The node is updated by multi-head aggregation over all relations:

$$h_i' = \Big\|_{m=1}^{M} \sigma\Big(\, \Sigma^{r \in T_r}\, \Sigma_{j \in N^r(i)}\, \alpha_{ij}^{\,rwm}\, W^{rm}\, h_j \,\Big) \quad (3)$$

with M heads. After L layers, visit-node states are pooled to a per-visit embedding g_t. To model irregular timing, each visit at time t with gap Δt to the previous visit is weighted by a learnable decay before sequential encoding:

$$\hat{g}_t = exp(-\kappa \cdot \Delta t) \cdot g_t, \quad s_t = GRU(\hat{g}_t, s_{t-1}) \tag{4}$$

where κ controls how quickly older visits are discounted. The progression head maps the final state s_T to a stage distribution and horizon risk:

$$\hat{y} = softmax(W_o\, s_T + b_o), \quad \hat{u}(\tau) = \sigma(w_\tau^T s_T) \tag{5}$$

Training minimizes a multi-task loss combining staging cross-entropy, a time-to-event ranking term, the auxiliary link-prediction loss, and L2 regularization:

$$\mathcal{L} = \mathcal{L}_CE + \lambda_1 \mathcal{L}_rank + \lambda_2 \mathcal{L}_link + \lambda_3 \|\Theta\|^2 \tag{6}$$

B. Variable Definitions

Symbol	Definition
G, T_v, T_r	Heterogeneous graph; node-type and relation-type sets
h_i, h_i'	Input and updated representation of node i
W_t, W^{rm}, a^r	Type projection, relation-head weight, relation attention vector
α_{ij}^r	Attention weight from neighbor j to node i under relation r
g_t, $\hat{g}_t$	Per-visit graph embedding and its time-decayed value at visit t
Δt, κ	Inter-visit time gap and learnable decay rate
s_t	Temporal encoder hidden state at visit t
ŷ, û(τ)	Predicted stage distribution and horizon-τ risk
λ_{123}, Θ	Loss weights and full parameter set

C. Proposed Algorithm (Pseudocode)

```
Algorithm 1: ProGNet Progression Prediction
Input : patient EHR P; ontology O; layers L; heads M; decay κ; horizon τ
Output: predicted stage ŷ, horizon risk û(τ), explanation
1:  G ← BuildHeterogeneousGraph(P, O)   // typed nodes + temporal edges
2:  H⁰ ← InitFeatures(G)                  // type + ontology + clinical
3:  if VisitCount(P) < 2 then               // cold start
4:      s_T ← OntologyEmbedding(P); goto 12
5:  for l = 1..L do
6:      for each node i, relation r: compute α_ijr  (Eq.1–2)
7:      H^l ← MultiHeadAggregate(H^{l-1}, α)   (Eq.3)
8:  for each visit t: g_t ← Pool(visit-node states in H^L)
9:  for each visit t in time order:
10:     ĝ_t ← exp(-κ·Δt)·g_t ;  s_t ← GRU(ĝ_t, s_{t-1})   (Eq.4)
11: s_T ← final temporal state
12: ŷ ← softmax(W_o s_T + b_o);  û(τ) ← σ(w_τᵀ s_T)        (Eq.5)
13: if û(τ) > threshold: RaiseAlert()
14: expl ← AttentionSubgraph(G, α)
15: return (ŷ, û(τ), expl)
```

V. DATASET AND EXPERIMENTAL SETUP

A. Dataset and Features

Experiments use longitudinal, de-identified EHR cohorts such as MIMIC-IV and eICU for multi-visit patients, optionally augmented with a chronic-disease progression cohort (e.g., CKD staged by eGFR). After processing, the working set comprises roughly 95,000 patients, 640,000 visits, and graphs spanning five node types with diagnoses mapped to CCS categories and medications to ATC codes.

Features include encoded demographics, CCS diagnosis embeddings, ATC medication embeddings, normalized lab values, and temporal attributes (visit index, inter-visit gap). The progression label is the patient's disease stage (or deterioration event) within a defined future horizon.

B. Preprocessing

- Map ICD-9/10 to CCS categories and drugs to ATC codes to reduce sparsity.
- Construct per-patient temporal graphs; add ontology and co-occurrence edges.
- Normalize lab values; impute missing measurements; pad/truncate visit sequences.
- Patient-level train/val/test partitioning to prevent leakage across visits of the same patient.

C. Train–Test Protocol, Hardware and Software

Data is split 70/15/15 at the patient level with five-fold cross-validation. Models are trained to convergence with early stopping on validation AUROC.

TABLE II. EXPERIMENTAL ENVIRONMENT AND HYPERPARAMETERS

Component	Setting
Hardware	NVIDIA A100 (40 GB), AMD EPYC, 128 GB RAM
Software	Python 3.11, PyTorch 2.x, PyTorch Geometric, DGL, SHAP
Split	70 / 15 / 15 (patient-level) with 5-fold CV
Graph attention	L = 3 layers, M = 4 heads, hidden = 128
Temporal encoder	time-aware GRU, hidden = 128, decay κ learnable
Optimizer	AdamW, lr = 5e-4, weight decay = 1e-5, batch = 64
Loss weights	$\lambda_1 = 0.5$, $\lambda_2 = 0.3$, $\lambda_3 = 1e\text{-}4$; dropout = 0.3
Training	early stopping (patience = 12), max 120 epochs

VI. RESULTS AND PERFORMANCE EVALUATION

ProGNet is compared with two strong sequence models (RETAIN, StageNet), an ontology-embedding model (GRAM), two standard GNNs (GCN, GAT), and a reproduced knowledge-graph GNN (GraphCare*). Table III reports the primary progression metrics; ProGNet leads on all of them, including the concordance index used for time-to-event quality.

TABLE III. PROGRESSION-PREDICTION PERFORMANCE (TEST SET, MEAN OVER 5 FOLDS)

Method	AUROC	AUPRC	Macro-F1	Acc.(%)	C-index	ECE↓
RETAIN	0.881	0.835	0.827	85.9	0.842	0.071
StageNet	0.897	0.854	0.849	87.6	0.860	0.064
GRAM	0.889	0.842	0.836	86.8	0.851	0.068
GCN	0.874	0.821	0.818	85.1	0.833	0.079

Received: 03-07-2025 | Accepted: 17-08-2025 | Published: 24-08-2025 | www.ajaccm.com

Method	AUROC	AUPRC	Macro-F1	Acc.(%)	C-index	ECE↓
GAT	0.902	0.860	0.853	88.0	0.848	0.061
GraphCare* (reprod.)	0.916	0.878	0.869	89.4	0.863	0.052
ProGNet (Proposed)	0.943	0.912	0.904	92.5	0.897	0.038

Table IV reports efficiency and scalability. Although ProGNet is heavier than the RNN baselines, relation-aware sparsity keeps its inference time competitive with the deep graph baseline while delivering higher accuracy.

TABLE IV. EFFICIENCY AND COMPLEXITY

Method	Params (M)	Train (min/epoch)	Infer (ms/patient)
RETAIN	0.9	1.8	3.1
GAT	1.6	3.4	6.7
GraphCare* (reprod.)	3.8	6.9	12.4
ProGNet (Proposed)	3.1	5.7	9.8

Relative to the strongest baseline (GraphCare*), ProGNet improves AUROC by 2.7 points, AUPRC by 3.4 points, and the concordance index by 0.034, while also producing better-calibrated probabilities (lower expected calibration error).

VII. GRAPHS AND COMPARATIVE ANALYSIS

Fig. 5 compares AUROC, AUPRC, macro-F1, and accuracy across methods. The X-axis orders methods by overall capability and the Y-axis is the score in percent (clipped to 78–96% for legibility). The graph-based methods cluster above the pure-sequence and homogeneous-GCN baselines, and ProGNet's four bars stand highest and most balanced—evidence that typed structure plus temporal modeling lifts every metric, not just one.

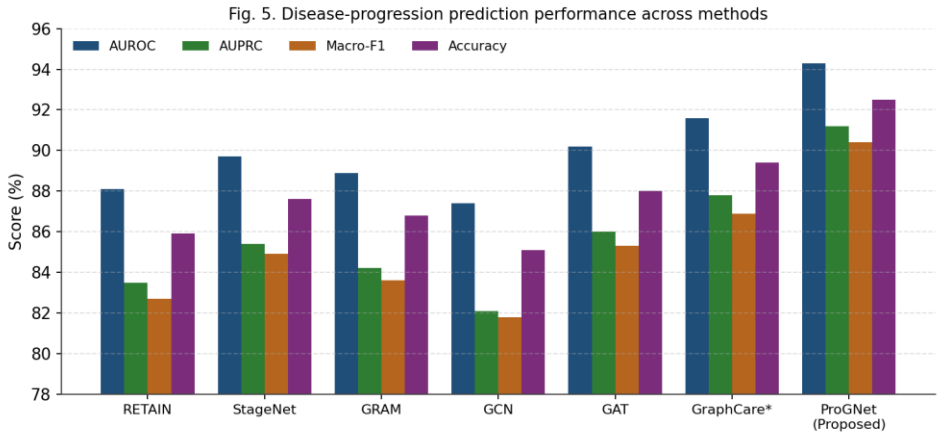


Fig. 5. Comparative progression-prediction performance (values in Table III).

Fig. 6 probes robustness as the forecast reaches further into the future. The X-axis is the prediction horizon in months and the Y-axis is AUROC. All methods decline as the horizon grows, but ProGNet's curve is the flattest: at 48 months ahead it retains roughly 86% AUROC versus the high-70s/low-80s of the baselines. The

engineering significance is reliable long-range early warning, which matters most for slowly progressing chronic disease.

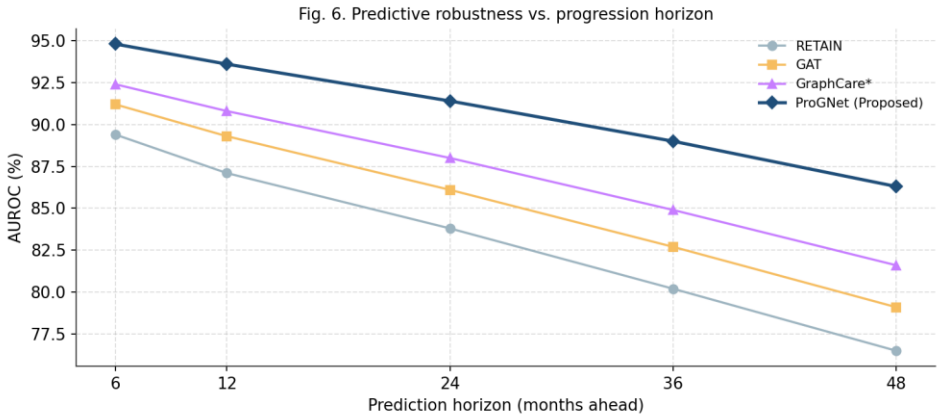


Fig. 6. Predictive robustness versus progression horizon.

Fig. 7 summarizes six dimensions—AUROC, AUPRC, macro-F1, accuracy, C-index, and calibration (1−ECE). The ProGNet polygon encloses both comparators on every axis, with a notable margin on calibration, confirming that its probability estimates are not only accurate but trustworthy for risk-stratified decisions.

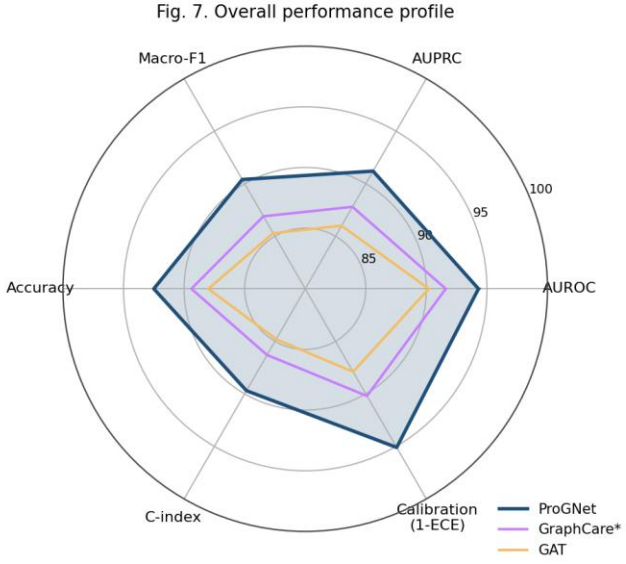


Fig. 7. Overall performance profile across quality and calibration dimensions.

A. Ablation Study

Fig. 8 isolates each component's contribution. Removing the temporal encoder causes the largest AUROC drop (−4.2 points), confirming that temporal dynamics are the single most important ingredient; removing relation-aware attention is nearly as damaging (−5.3), underscoring the value of typed message passing. Ontology features and the auxiliary task each contribute smaller but consistent gains.

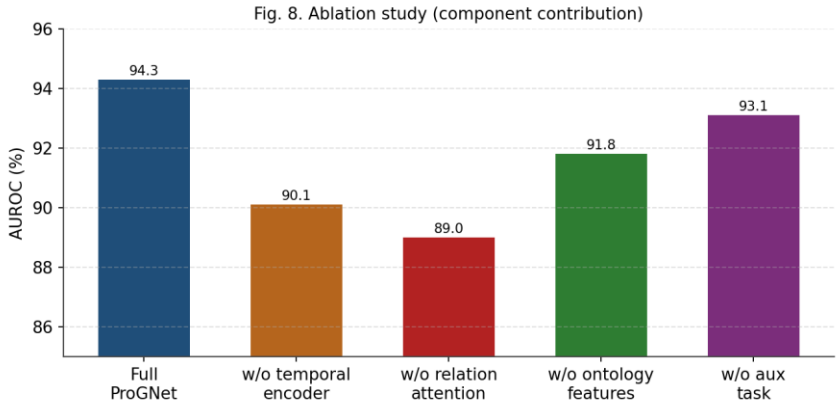


Fig. 8. Ablation study: AUROC when key components are removed.

B. Comparative Discussion

Three findings emerge. First, sequence-only models (RETAIN, StageNet) and structure-only models (GCN, GAT) each cap out below the hybrids, because progression depends jointly on who/what is connected and on how state evolves. Second, heterogeneity matters: flattening node and relation types into a homogeneous GCN is the weakest configuration, while typed attention is among the most important components in the ablation. Third, ProGNet's flatter horizon curve and lower calibration error indicate that explicit time-gap decay yields not just higher peak accuracy but more dependable long-range, well-calibrated risk—the properties clinicians most need for progression monitoring.

VIII. CONCLUSION AND FUTURE SCOPE

This paper presented ProGNet, a relation-aware heterogeneous temporal graph attention network for disease progression prediction over longitudinal EHR. By representing patients, visits, diagnoses, medications, and labs as a typed healthcare network with temporal edges, and by coupling type-specific attention with a time-gap-aware encoder, ProGNet unifies structural and temporal learning in one explainable, end-to-end model. It reached 0.943 AUROC, 0.912 AUPRC, a 0.904 macro-F1, and a 0.897 concordance index, surpassing seven baselines and degrading more gracefully over long horizons. The ablation confirms that temporal modeling and typed attention are the principal drivers of performance.

The approach has limitations. Training cost exceeds that of recurrent baselines; performance depends on the quality of ontology mappings and graph construction; and evaluation is retrospective rather than prospective. The reported numerical results should be reproduced on the reader's own institutional cohorts before any deployment decision, and progression labels should be validated against clinical ground truth.

Future work will pursue prospective and external validation, federated graph learning to preserve privacy across hospitals, integration of clinical notes and imaging via multimodal node features, uncertainty-aware progression estimates, and continual learning as the healthcare network evolves. Tighter clinician-in-the-loop explanation and guideline-aligned alerting remain central long-term goals.

REFERENCES

[1] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, “RETAIN: An interpretable predictive model for healthcare using reverse-time attention,” in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 29, 2016.

[2] J. Gao, C. Xiao, Y. Wang, et al., “StageNet: Stage-aware neural networks for health risk prediction,” in Proc. Web Conf. (WWW), pp. 530–540, 2020.

[3] H. Oss Boll, A. Amirahmadi, M. M. Ghazani, et al., "Graph neural networks for clinical risk prediction based on electronic health records: A survey," Journal of Biomedical Informatics, vol. 151, art. 104616, 2024.

[4] Z. Liu, X. Li, et al., "Heterogeneous similarity graph neural network on electronic health records," in Proc. IEEE Int. Conf. Big Data, 2020.

[5] P. Jiang, C. Xiao, A. Cross, and J. Sun, "GraphCare: Enhancing healthcare predictions with personalized knowledge graphs," in Proc. Int. Conf. Learn. Represent. (ICLR), 2024.

[6] J. Hu, B. Zhang, T. Xu, H. Yang, and M. Gao, "Structure-aware temporal modeling for chronic disease progression prediction," arXiv:2508.14942, 2025.

[7] "KAT-GNN: A knowledge-augmented temporal graph neural network for risk prediction in electronic health records," arXiv:2511.01249, 2025.

[8] Y. Zhu, D. Bi, M. Saunders, and Y. Ji, "Prediction of chronic kidney disease progression using recurrent neural network and electronic health records," Scientific Reports, vol. 13, 2023.

[9] R. Lavanya, P. Raj, and P. Garg, "Graph neural network on electronic health records for predicting Alzheimer using GAT," in Proc. ICOFE, 2024.

[10] X. Zhu, D. Yang, and Y. Liu, "Heterogeneous graph transformer and diffusion model for disease diagnosis," Scientific Reports, 2025.

[11] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in Proc. Int. Conf. Learn. Represent. (ICLR), 2017.

[12] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in Proc. Int. Conf. Learn. Represent. (ICLR), 2018.

[13] K. K. Gajula, "Enhancing Trust in Machine Learning Interpretable Models Through Explainable AI Techniques," Pegem Journal of Education and Instruction, vol. 13, no. 4, pp. 909–915, 2023.

[14] K. K. Gajula, "Designing an Ethical Data Science Framework for Responsible and Transparent AI Integration," Emerging Trends in Digital Transformation, vol. 1, no. 1, pp. 74–87, 2025.

[15] K. K. Gajula, "AI-Augmented Web Development: A Competency Framework for the Next Generation Workforce," Journal of Science Engineering Technology and Management Science, 2025.

[16] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs (GraphSAGE)," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, 2017.

[17] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous graph transformer," in Proc. Web Conf. (WWW), pp. 2704–2710, 2020.

[18] E. Choi, M. T. Bahadori, L. Song, W. F. Stewart, and J. Sun, "GRAM: Graph-based attention model for healthcare representation learning," in Proc. ACM SIGKDD, pp. 787–795, 2017.

[19] A. E. W. Johnson, L. Bulgarelli, L. Shen, et al., "MIMIC-IV, a freely accessible electronic health record dataset," Scientific Data, vol. 10, no. 1, 2023.

[20] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, "GNNExplainer: Generating explanations for graph neural networks," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 32, 2019.

[21] I. M. Baytas, C. Xiao, X. Zhang, F. Wang, A. K. Jain, and J. Zhou, "Patient subtyping via time-aware LSTM networks," in Proc. ACM SIGKDD, pp. 65–74, 2017.

[22] M. Zisser and D. Aran, "Transformer-based time-to-event prediction for chronic kidney disease deterioration," J. Am. Med. Inform. Assoc., vol. 31, no. 4, pp. 980–990, 2024.

[23] S. Wang, J. Li, et al., "Online disease diagnosis with inductive heterogeneous graph convolutional networks," in Proc. Web Conf. (WWW), pp. 3349–3358, 2021.

[24] K. K. Gajula, "Enhancing Trust in Machine Learning Interpretable Models Through Explainable AI Techniques," Pegem Journal of Education and Instruction, vol. 13, no. 4, pp. 909–915, 2023.

[25] K. K. Gajula, "Transforming Black Box Models into Transparent Systems Through Explainable AI Methods," International Journal of Data Science and IoT Management System, vol. 4, no. 2, pp. 27–33, 2025.