

Trajectory-Aware Clinical Risk Prediction via Severity-Grounded Knowledge Graphs and Retrieval-Augmented Generation

Kyunghoon Jeon
Hanyang University
Seoul, Republic of Korea
kyunghoon123@hanyang.ac.kr

Woohwan Jung
Korea University
Seoul, Republic of Korea
woohwan@korea.ac.kr

Youmin Ko
Hanyang University
Seoul, Republic of Korea
youminkk0213@hanyang.ac.kr

Hyunjoon Kim*
Hanyang University
Seoul, Republic of Korea
hyunjoonkim@hanyang.ac.kr

Abstract

While Electronic Health Records (EHRs) offer a wealth of clinical data, effectively augmenting a patient's records with heterogeneous external knowledge to predict the patient's clinical risk remains a significant challenge. Existing methods fail to capture disease severity, treatment responses, and nuanced clinical progression, due to data sparsity and the underutilization of unstructured clinical notes. To address these challenges, we propose **TRACER** (*a trajectory-aware and clinically grounded prediction framework*) that (1) constructs a medical knowledge graph enriched with severity information from medical literature, (2) retrieves clinically relevant, severity-weighted paths of a patient's progression from the knowledge graph, (3) extracts clinically relevant events from unstructured clinical notes, and (4) augments patient context with similar peer cases. Experiments on the MIMIC-III and MIMIC-IV datasets demonstrate large gains over state-of-the-art baselines, with up to 28.5% increase in Macro F1 score for the mortality prediction task, and 19.7% increase for the readmission prediction task. Beyond performance, TRACER yields transparent, path-level explanations, enabling clinicians to inspect the literature-backed logic and specific textual evidence driving each risk estimate.

CCS Concepts

• **Applied computing** → **Health care information systems**;
• **Information systems** → *Information retrieval*; • **Computing methodologies** → **Knowledge representation and reasoning**.

Keywords

Electronic Health Records; Healthcare Outcome Prediction; Knowledge Graph; Severity Scoring; Retrieval-Augmented Generation

ACM Reference Format:

Kyunghoon Jeon, Youmin Ko, Woohwan Jung, and Hyunjoon Kim. 2026. Trajectory-Aware Clinical Risk Prediction via Severity-Grounded Knowledge Graphs and Retrieval-Augmented Generation. In *Proceedings of the*

32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26), August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770855.3818040>

Resource Availability:

The source code of this paper has been made publicly available at <https://github.com/KyunghoonJeon/TRACER.git>.

1 Introduction

Clinical risk prediction aims to estimate the clinical risk of adverse outcomes, including mortality [28, 49] and readmission [1, 12], and can support early interventions that help prevent these critical events. To this end, recent methods for clinical prediction rely on Electronic Health Records (EHRs), which contain patients' time-stamped medical histories associated with medical concepts (e.g., diagnoses, procedures, medications), their demographic data, and clinical notes. Building upon EHRs, recent studies have leveraged knowledge graphs (KGs) [7, 41, 47] and have retrieved relevant medical information from external literature [44] such as PubMed [4]. Notably, several approaches [26, 37] construct a patient-specific subgraph of the KG typically associated with medical concepts from a target patient's visit sequence.

Despite the recent advances in clinical risk prediction, these methods still have several limitations. First, they often struggle to cope with the sparsity and irregularity of patient visit histories in real-world EHRs. Most patients have only a limited number of recorded visits, often separated by heterogeneous and prolonged intervals.

Under such conditions, existing sequential models may fail to capture meaningful temporal dynamics or identify critical turning points in disease progression. Although KG-based approaches [15, 16] attempt to mitigate this issue by incorporating external medical knowledge, they selectively individually extract the medical knowledge related to each of the medical concepts included in a patient's EHR, typically treat all medical concepts equally, ignoring the fact that different diagnoses may have different levels of clinical importance. As a result, these methods fail to capture severity of concepts and highly plausible disease progressions across visits, thereby being prone to overlook clinically critical "red-flag" events. This limitation is exacerbated by the extreme sparsity of real-world data. As shown in Figure 1, 68.6% of patients in the MIMIC-III dataset have only a single visit, and the vast majority

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. KDD '26, Jeju Island, Republic of Korea
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2259-2/2026/08
<https://doi.org/10.1145/3770855.3818040>




Figure 1: Macro F1 and sensitivity across visit sequence lengths for the mortality prediction task, with the percentage of patients by number of visits.

(95.7%) have fewer than five recorded visits. Under such sparse conditions, state-of-the-art baselines such as GraphCare, DearLLM, and KARE exhibit significant performance degradation. For the mortality prediction task over patients with limited history with ≤ 5 visits, the Macro F1 scores consistently remain below 0.7 and sensitivity fails to exceed 0.5 except KARE for five visits¹. Second, they underutilize the rich and nuanced information contained in clinical notes. Clinical narratives often encode high-value signals such as physicians’ assessments, suspected complications, treatment responses, and subtle changes in patient condition. However, prior methods either discard these textual sources altogether or employ only a limited set of extracted medical concepts as the patient features. Such simplifications strip away clinically salient context and narrative evidence, which can be crucial for identifying early warning signs and for explaining adverse clinical outcomes.

These limitations provide important insights into three potential solutions: (i) augmenting a patient’s sparse visit sequence with the connected knowledge highly associated with that sequence from medical KGs and the literature, to reveal underlying health progression of her/his health status; and (ii) leveraging the nuanced textual context in clinical notes to capture subjective patients’ complains and reactions often underexplored by the existing methods; and (iii) relying solely on a patient’s own history often fails to leverage the clinical experience embedded in historically similar patient cohorts, particularly when personal data is scarce.

Guided by these insights, we propose TRACER, a new retrieval-augmented generation (RAG) framework for trajectory-aware and clinically grounded clinical risk prediction, which explicitly accounts for disease severity, personal narrative evidence, and cohort-level experience. First, to address sparse visit histories and the uniform treatment of diagnoses, TRACER constructs a novel medical knowledge graph called SMKG, grounded on the severity by assigning clinically meaningful scores to diagnoses based on curated medical literature and clinical guidelines. Severity grounding enables the model to distinguish clinically critical diagnoses and red-flag events from less consequential conditions. Given a patient’s longitudinal visit sequence, TRACER retrieves severity-aware medical trajectories which are sequential paths connecting concepts across adjacent visits that capture clinically plausible disease progressions even under sparse observations. Moreover, to move beyond reliance

on an individual patient’s history, TRACER retrieves clinically similar patient cases based on recent visit patterns and trajectories, allowing an LLM to leverage outcome patterns and progression dynamics observed in historically similar cohorts. Second, to fully exploit textual evidence, TRACER enriches each retrieved trajectory with salient evidence from clinical notes, capturing nuanced signals such as clinician assessments, suspected complications, treatment responses, and patient-reported symptoms that are often omitted or oversimplified in structured EHR codes. By jointly integrating severity-aware KG trajectories, clinical textual evidence, and similar patient retrieval in the RAG framework, TRACER not only improves predictive performance under real-world EHR sparsity but also provides path-level, clinically interpretable explanations.

Our main contributions are summarized as follows:

- We propose the severity-grounded medical knowledge graph (SMKG) by assigning fine-grained and diagnosis-level severity scores derived from medical literature.
- We address the data sparsity problem by augmenting (a) the target patient’s sparse data with patient-specific and severity-aware progression paths across visits from SMKG and (b) clinically similar patient cases, both of which enables the LLM reasoning based on nuanced patient status changes under data sparsity.
- We propose to leverage salient and nuanced textual evidence in clinical notes, enabling robust modeling of red-flag events and personal experiences.
- We demonstrate consistent and substantial performance gains on the MIMIC-III and MIMIC-IV datasets for two tasks, i.e., mortality prediction and readmission prediction, together with path-level and clinically grounded explanations that enhance interpretability and trust.

2 Related Work

Early methods [2, 5, 6, 10, 27, 30, 31, 34, 39] for clinical risk prediction primarily focus on modeling the temporal dynamics of patient visit sequence without external knowledge. In contrast, recent approaches focus on two directions: (i) constructing medical KGs from EHR data and extracting subgraphs from external KGs based on patient visit information, (ii) retrieving textual information from external knowledge sources to complement EHR data for clinical risk prediction.

Graph-based Clinical Risk Prediction. To capture the complex interactions among patient statuses and the relationships between diagnoses and treatments, several methods utilize KGs. GRAM [7] and KAME [29] construct medical KGs with attention to enrich patient embeddings. GRASP [52] clusters similar patients into groups and models the relationships between group centroids as a graph. KerPrint [50] builds time-aware personalized KG to model disease progression. GraphCare [15] builds personalized KGs by integrating each patient’s diagnoses, procedures, and medications. However, they extract only 1-hop neighbors of patient-relevant medical concepts in the KGs, and ignore the severity of the medical concepts, thus being unable to capture high-risk events and overlook hidden clinical progressions between visits.

Text-based Clinical Risk Prediction. Previous text-based methods [9, 51] retrieve clinically relevant information from structured

¹Analogous trends are observed in the readmission prediction task, with details provided in Appendix A. The decrease in Macro F1 and increase in sensitivity across visit-count groups reflect differences in outcome prevalence across subgroups, and the same trend is observed across all compared models.

EHR data. Recently, LLMs have been applied to healthcare tasks due to their advanced reasoning capabilities. KARE [16] integrates EHRs with external knowledge via tailored prompts to generate prediction evidence. DearLLM [48] quantifies associations between medical concepts by measuring how likely an LLM is to predict one concept given another, using patient-specific prompts. CPLLM [38] fine-tunes LLMs on linearized EHR data. Additionally, several approaches leverage LLMs for clinical text comprehension, such as retrieving and summarizing relevant documents [46] or extracting medical concepts from clinical notes [53]. However, these methods either neglect clinical notes or fail to fully utilize their rich textual information. They also overlook the severity of diagnoses, resulting in missed nuanced clinical knowledge.

3 Preliminary

3.1 Electronic Health Records

A medical concept $c \in C$ refers to a diagnosis, a procedure, or a medication, where C is the set of all medical concepts, e.g., hypertension (diagnosis), appendectomy (procedure), and metformin (medication). Electronic Health Records (EHRs) comprise visit records, demographic data, and clinical textual records of patients. For each patient, her/his visit records are represented as the temporally ordered visit sequence $[x_1, x_2, \dots, x_t]$, where each visit $x_i \subseteq C$ consists of the diagnoses, procedures, and medications provided at the patient's i -th visit. The demographic d_i contains the patient's attributes available at the i -th visit, such as ethnicity, age and gender. We denote a set of clinical notes for the patient at visit i as $\mathcal{R}_i = \{r_{i,1}, \dots, r_{i,m_i}\}$, where each clinical note represent a progress note, discharge summary, or physician note, and m_i denotes the number of clinical notes in $\mathcal{R}_i$. The patient's all clinical textual records are represented as $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_t\}$. More information of EHRs is detailed in the extended version of this work [14].

3.2 Problem Formulation

For clinical risk prediction tasks such as mortality prediction and readmission prediction, our goal is to predict specific clinical outcomes based on a patient's EHRs up to that patient's latest visit at T . The patient's EHRs are represented as her/his historical visit sequence $[x_1, \dots, x_T]$, demographic d_T , and the set $\mathcal{R}$ of clinical notes. The predicted outcomes are defined as follows.

Mortality Prediction. Mortality prediction aims to predict a patient's survival status $y \in \{0, 1\}$, where 0 indicates the patient will be alive, and 1 indicates the patient will be deceased.

Readmission Prediction. Readmission prediction aims to determine a binary outcome $y \in \{0, 1\}$ that denotes whether the patient will be readmitted within σ days after discharge from the latest historical visit x_T , with 0 indicating no readmission and 1 indicating readmission. Following previous work [15, 16], we set $\sigma = 15$ days.

4 Methodology

4.1 Overview of the TRACER Framework

Figure 2 demonstrates the TRACER framework that consists of three stages: (1) patient medical profile (PMP) retrieval, (2) similar patient retrieval, and (3) clinical risk prediction. In stage (1), given a target patient's EHR and our medical KG, TRACER retrieves the




Figure 2: Overview of TRACER




Figure 3: Construction of the severity-weighted medical knowledge graph (SMKG)

patient's medical profile (PMP), which consists of compact patient-relevant trajectories, relevant passages from the patient's clinical notes, and demographics. In stage (2), we retrieve the trajectories and demographic attributes of top- m patients most similar to the target patient. In stage (3), the LLM generates binary predictions for mortality or readmission, based on the patient's medical profiles, and the top- m similar patients' context. Next, we will introduce our medical KG construction method and describe each of the three main stages.

4.2 Severity-weighted Medical KG Construction

Despite the importance of considering the severity of diagnoses, existing studies utilizing medical KGs [6, 31, 52] treat all diagnoses equally without reflecting differences in their severity. To address this limitation, we propose a retrieval-augmented severity generation method that quantifies the clinical severity of diagnostic codes, as shown in Figure 3.

First, we construct a unified medical KG by integrating three data sources: the existing medical KG such as UMLS [3], medical documents in PubMed [4], and LLM-generated triples to capture comprehensive relationships between medical concepts, which have been utilized by KARE [16]. Subsequently, for each diagnostic concept in the unified medical KG, the passage retriever searches all abstracts in PubMed and all summaries in Wikipedia, respectively, for top- l passages closely related to that diagnosis. The LLM takes the concatenation of the retrieved passages as input and generates a severity score for that diagnosis. The score ranges from 1 to 20,




Figure 4: Architecture of the profile retriever

where higher values indicate greater clinical severity². We then define the severity-weighted medical KG (SMKG) as a KG that shares the same node and edge sets as the unified medical KG, where each diagnosis node is assigned the severity score of that diagnosis, and the other types of nodes are assigned a weight of 0³.

In subsequent stages, these severity scores enable a more granular capture of a patient’s clinical progression patterns by quantitatively reflecting the clinical severity of each diagnosis, in contrast to existing approaches. We provide an analysis of scoring granularity in Appendix C. Additionally, we further validate the clinical relevance of the proposed severity scores by analyzing their correlation with mortality outcomes in Appendix D.

This LLM-based severity scoring is motivated by recent work showing that LLMs can perform clinically meaningful risk stratification and severity-related assessment [8, 33]. Since no universal diagnosis-level severity scale exists across heterogeneous diseases, we additionally assessed the generated scores using a four-level disease-severity framework [23]. The mean absolute deviation was 0.06 across three runs, corresponding to a 1.6% error. This indicates that the generated severity scores are reasonably aligned with coarse clinical severity categories. We use severity not as a standalone prognostic label, but as a retrieval prior that is further contextualized by patient-specific trajectories and clinical notes.

4.3 Patient Medical Profile Retrieval

Most patients’ EHR data are sparse. Their visit records typically contain one to five visits, and they are often spaced far apart chronologically. As shown in Figure 4, we address this issue by augmenting the target patient’s sparse EHR with her/his structured and unstructured medical history which is highly relevant to the patient and our tasks, referred to as the patient medical profile (PMP). PMP includes (1) the task-relevant progression of the patient’s diseases and her/his responses to treatments, (2) the task-relevant descriptions of her/his health status, and (3) the demographic of the patient. For each patient, we retrieve two types of medical profiles focusing on

two complimentary clinical factors. The risk factors are associated with adverse outcomes, such as progressive organ dysfunction and high-risk comorbidities; the protective factors are indicative of clinical stability or improvement, such as recent treatment response and stable vital signs.

4.3.1 Trajectory Retrieval. We hypothesize that the evolution of a patient’s health status can be modeled as a path through the underlying knowledge graph. Formally, we define a trajectory p as a path in the SMKG that connects medical concepts across consecutive visits, following temporal order. For every pair of adjacent visits x_{t-1} and x_t in the visit sequence where $t \in \{2, \dots, T\}$, we consider all cross-visit concept pairs $(c_i, c_j) \in x_{t-1} \times x_t$. For each pair, we extract all shortest paths in the SMKG up to a maximum length of $h = 3$ hops to ensure both clinical relevance and computational tractability. Let $\mathcal{P}_t$ denote the set of all such visit-to-visit paths between x_{t-1} and x_t . The complete set of candidate trajectories is then formed by the Cartesian product $\mathcal{P}_2 \times \dots \times \mathcal{P}_T$, where each candidate trajectory corresponds to a unique sequential concatenation of paths selected from $\mathcal{P}_2, \dots, \mathcal{P}_T$ ⁴. We compute the embedding $\mathbf{e}_p$ of each candidate trajectory using the pretrained biomedical encoder BioClinicalBERT [24].

We preserve these path-level trajectories rather than merging them into a single unified patient graph; this design allows them to be independently ranked, diversified, repacked, and directly verbalized as explicit, interpretable evidence for the LLM.

To capture the full spectrum of a patient’s medical profile, we perform a dual-retrieval process targeting both risk-inducing and protective trajectories. We define risk factors as clinical evidence associated with adverse outcomes, such as progressive organ dysfunction, severe comorbidities, or deteriorating physiological status. In contrast, protective factors refer to evidence indicating clinical stability or recovery, such as treatment response, stable vital signs, or improving symptoms [13, 22]. Based on these definitions, we construct two distinct textual queries: a risk query q_r and a protective query q_o . Each query is composed of: (1) the specific prediction task description and (2) the patient’s demographics d_T . For every

²This fine-grained severity scale 1–20 enables the model to distinguish subtle yet clinically meaningful differences among diagnoses, preventing high-risk conditions from being diluted by low-impact or incidental events [32].

³Other medical concepts such as medications or procedures are excluded from severity scoring, as clinical severity is typically defined for disease conditions rather than interventions or treatments.

⁴Single-visit patients are excluded from the cohort; for retained patients, the final visit serves as the prediction target, while preceding visits constitute the historical context.

candidate trajectory p , we compute two relevance scores ψ_r and ψ_o for risk and protective queries, respectively:

$$\begin{aligned}\psi_r(q_r, p) &= \alpha \cos(\mathbf{e}_{q_r}, \mathbf{e}_p) + (1 - \alpha)s(p) \\ \psi_o(q_o, p) &= \beta \cos(\mathbf{e}_{q_o}, \mathbf{e}_p) + (1 - \beta)(1 - s(p))\end{aligned}$$

where $\mathbf{e}_{q_r}$ and $\mathbf{e}_{q_o}$ denote the embeddings of the respective query, and the severity $s(p)$ of trajectory p is calculated as the average of the normalized severity score (ranging $[0, 1]$) of every diagnosis node contained in p . Risk and protective aspects are inherently complementary: an increase in one severity implies a decrease in the other. By using $1 - s(p)$ for ψ_o , we maintain interpretive consistency and joint expressiveness. Hyperparameters $\alpha, \beta \in [0, 1]$ weigh the relative importance of the semantic relevance to the queries and the severity of the trajectory. Finally, we retrieve the top $\rho\%$ of trajectories for each aspect based on these scores. This linear score is not used as a standalone risk estimator; rather, it serves as an interpretable retrieval prior that prioritizes trajectories that are both clinically severe and relevant to the target prediction task.

4.3.2 Trajectory Refinement. While coarse first-pass retrieval in the previous stage provides query-relevant results, it often yields candidates that lack precise relevance and exhibit redundancy. Such redundancy may not only reduce the diversity of retrieved evidence but also obscure clinically meaningful progression patterns that are critical for downstream prediction and interpretation. In the trajectory refinement stage, we thus rerank the retrieved trajectories, and filter out the redundant trajectories among them to achieve both the query-trajectory relevance and diversity. Specifically, we initialize an empty pool S_0 of trajectories. Next, we iteratively refine them based on maximum marginal relevance (MMR) to balance between the relevance of the trajectories to the aforementioned risk (spec. protective query) and their novelty with respect to the already selected set S_{i-1} . At each iteration i , we compute the below gain score for each of the retrieved trajectories not included in S_{i-1} , and add the trajectory p^* with the highest gain score to the pool, i.e., $S_i = S_{i-1} \cup \{p^*\}$.

$$g(p_i) = \lambda_{MMR} \cos(\mathbf{e}_{q_i}, \mathbf{e}_{p_i}) - (1 - \lambda_{MMR}) \max_{j \in S_{i-1}} \cos(\mathbf{e}_{p_i}, \mathbf{e}_{p_j})$$

where hyperparameter λ_{MMR} is a weighting factor. Beyond semantic similarity, we further assess whether the selected trajectory set collectively entails the clinical conditions and diagnoses closely associated with the target patient. We calculate the entailment score η_i of S_i by using the pretrained NLI model SimCSE-Roberta [11]:

$$\eta_i = NLI(\{p|p \in S_i\}, q)$$

where $NLI(\cdot, \cdot)$ takes as input the text for all trajectories $p \in S_i$ and the query q , i.e., either q_r or q_o . The iteration terminates when either $\eta_i \geq \lambda_{max}$ or $\eta_i - \eta_{i-1} \leq \lambda_{min}$ is satisfied where thresholds λ_{max} and λ_{min} are hyperparameters. This combination of MMR and NLI-based stopping preserves relevant, diverse, and sufficiently supportive trajectories, aligned with prior retrieval-augmented evidence selection paradigms [18]. We denote the trajectories remaining in the candidate pool after termination as Key Supporting Trajectories (KSTs). Finding the KSTs can be achieved solely by designing the prompt and adjusting the thresholds, without requiring additional model training. The KSTs compose context provided to the LLM that performs the prediction tasks.

4.3.3 Clinical Note Retrieval. Patient clinical notes are unstructured documents that comprehensively record various information including physicians' subjective assessments, detailed laboratory measurements, symptom descriptions, and the rationale behind prescriptions. Such unstructured narratives often contain early warning signs and nuanced clinical cues that are not explicitly captured by structured EHR codes. However, utilizing all clinical notes without filtering may introduce irrelevant information, which can degrade the performance of predictive models, e.g., LLMs.

To address this issue, we retrieve relevant passages from the set $\mathcal{R}$ of the patient's historical clinical records. Recall that $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_T\}$ contains all clinical textual records associated with the patient's historical visits. We construct the passage pool $\mathcal{B}$ by splitting every clinical note in $\mathcal{R}$ into passages. For each of risk query q_r and protective query q_o , we embed the query and each passage $b \in \mathcal{B}$ with a pretrained encoder, compute cosine similarities between the query and passage embeddings, and select the top- n passages most relevant to the query. The retrieved passages are then used as note-level evidence in the patient medical profile.

4.4 Retrieval-augmented Clinical Risk Prediction

Patients' EHRs typically consist of a small number of visits. However, the existing methods [10, 48] rely solely on a patient's own records, and embed the entire visit sequence into single vectors, which often introduce irrelevant or outdated information. They also [42, 54] do not take advantage of valuable information hidden in their unstructured clinical notes. To address these limitations, we augment a target patient's EHR with **three types of supporting evidence**: (A) the patient's risk and protective medical profiles, (B) the patient's demographics, and (C) the KSTs and demographics of her/his top- m similar patients. Specifically, the risk and protective medical profiles are produced by the profile retriever described in Section 4.3. For evidence (C), we first calculate the Jaccard similarity between the set of medical concepts in the target patient's recent visits and those from all other patients, and we then select the top- m similar patients based on the similarity. Next, we concatenate the healthcare task description, and evidence (A), (B) and (C) to form a single prompt. During this, we perform repacking on the KSTs and clinical notes, respectively, in evidence (A) so that the more relevant KST and clinical note come earlier and are thus located closer to the task description in the prompt. This repacking strategy allows the LLM to focus on more relevant information [40, 45]. Eventually, the LLM generates a binary outcome $y \in \{0, 1\}$.

By harmonizing (i) retrieval-augmentation with the severity-aware trajectories, the relevant clinical notes, and similar peer cases, (ii) dynamic context repacking, and (iii) a step-by-step reasoning process, our framework mitigates the "lost-in-the-middle" problem [25] inherent to long-context LLM inputs. It addresses both sparsity in time-series data and redundancy in unstructured notes. This retrieval-augmented, repacking-guided reasoning not only improves predictive performance but also generates transparent and interpretable chains of evidence that substantiate each risk estimate, thereby enhancing the clinical trustworthiness of the predictions. The impact of similar patients and the analysis of the repacking strategy are provided in Section 5.3 and Appendix B, respectively.

Table 1: Overall performance of two tasks on MIMIC-III. The numbers in parentheses indicate the proportion of patients with positive labels, i.e., 1s, for each task. The best and second best results are denoted in bold and underlined, respectively.

Models	Mortality Prediction (10.83%)				Readmission Prediction (21.6%)			
	Macro F1	Sensitivity	Specificity	Accuracy	Macro F1	Sensitivity	Specificity	Accuracy
GRASP	0.5176	0.0495	0.9963	0.8938	0.5262	0.1666	0.8897	0.7335
StageNet	0.5571	0.1881	0.9189	0.8390	0.5275	0.1666	<u>0.8923</u>	0.7355
MedRetriever	0.5622	0.2079	0.9122	0.8358	0.5746	0.6524	<u>0.6168</u>	0.6245
GraphCare	0.5963	0.2552	0.9577	0.9229	0.6364	0.6999	0.5964	0.6779
RAM-EHR	0.6165	0.2277	<u>0.9627</u>	0.8830	0.5678	0.7143	0.5761	0.6056
DearLLM	0.6374	0.2673	<u>0.9615</u>	0.8863	0.6390	0.7095	0.6194	0.7343
KARE	<u>0.6375</u>	<u>0.3711</u>	0.9593	<u>0.9267</u>	<u>0.7271</u>	<u>0.7501</u>	0.7112	<u>0.7451</u>
TRACER	0.8197	0.7128	0.9531	0.9275	0.8709	0.7809	0.9501	0.9135

Table 2: Overall performance of two tasks on MIMIC-IV. The numbers in parentheses indicate the proportion of patients with positive labels, i.e., 1s, for each task. The best and second best results are denoted in bold and underlined, respectively.

Models	Mortality Prediction (38.13%)				Readmission Prediction (48.84%)			
	Macro F1	Sensitivity	Specificity	Accuracy	Macro F1	Sensitivity	Specificity	Accuracy
GRASP	0.7514	0.5063	0.9688	0.7925	0.5876	0.5980	0.5766	0.5881
StageNet	0.7479	0.5494	0.9236	0.7809	0.6108	0.6594	0.5652	0.6112
MedRetriever	0.7727	0.5797	0.9392	0.8021	0.6237	0.6931	0.5595	0.6247
GraphCare	0.8356	0.6327	0.9491	0.8782	0.6713	0.6831	0.6219	0.6493
RAM-EHR	0.8223	0.6303	0.9579	0.8460	0.6374	0.6495	0.6257	0.6373
DearLLM	0.8130	0.6456	0.9516	0.8349	0.6870	0.7485	0.6295	0.6876
KARE	<u>0.8698</u>	<u>0.7498</u>	0.9771	<u>0.8893</u>	<u>0.7444</u>	<u>0.8582</u>	<u>0.6469</u>	<u>0.7256</u>
TRACER	0.8830	0.7722	<u>0.9688</u>	0.8938	0.7942	0.8752	0.7183	0.7949

Table 3: Ablation study on MIMIC-III. The variants are grouped by major evidence modules. Standard deviations over three runs are below 0.06 for all variants.

Module	Variant	Mortality		Readmission	
		Macro F1	Sensitivity	Macro F1	Sensitivity
Full model	TRACER	0.8197	0.7128	0.8709	0.7809
SMKG	w/o Sev.	0.7484	0.4950	0.8643	0.7571
Trajectory	w/o Risk.	0.7328	0.4356	0.8499	0.7095
	w/o Prot.	0.7318	0.5247	0.8507	0.7143
	w/o Traj.	0.6302	0.4653	0.8103	0.7000
Clinical notes	w/o Clin.	0.7757	0.5049	0.7981	0.6762
Similar patients	w/o Pat.	0.7527	0.5049	0.8050	0.7047
Combined	w/o PMP	0.6656	0.2475	0.7381	0.5285
	w/o (PMP+Pat.)	0.6424	0.2173	0.6999	0.4952

A qualitative case study illustrating trajectory-level evidence and reasoning process of TRACER is provided in Appendix L of the extended version [14].

5 Experiments

5.1 Experimental Setup

Datasets. We utilize two publicly available datasets, MIMIC-III [21] and MIMIC-IV [20]. The mortality rates are 10.83% and 38.13%, while the readmission rates are 21.6% and 48.84% in MIMIC-III and

MIMIC-IV, respectively. The extended version [14] details the label definition.

Evaluation Metrics. We assess performance using four metrics: Macro F1 score, sensitivity, specificity, and accuracy. A higher value indicates better performance.

Baseline Methods. We adopt state-of-the-art healthcare prediction models: StageNet [10], GRASP [52], MedRetriever [51], RAM-EHR [46], DearLLM [48], GraphCare [15], and KARE [16]⁵. We further compare TRACER with a recent EHR foundation model, ETHOS, and a prompt-only zero-shot LLM baseline to isolate the effect of retrieval augmentation. We showcase the details of the baseline implementation in [14].

Implementation Details. We conduct all experiments on a server equipped with an Intel Xeon(R) Gold 6226R CPU, a single NVIDIA RTX A6000 GPU (48 GB VRAM), and 256 GB of DDR4 RAM. For risk and protective trajectory retrieval, α and β are set to 0.8 and 0.2 respectively. λ_{MMR} in the trajectory refinement stage is set to 0.7, and the thresholds λ_{max} and λ_{min} for the entailment scores are 0.5 and 0.05 respectively. For clinical notes and similar patients retrieval, we set both n and m to 5. We leverage OpenAI GPT-4o-mini [35] to assign diagnostic severity scores and to perform retrieval-augmented clinical risk prediction. For fair comparison, we use GPT-4o-mini as the backbone LLM for all LLM-based baselines.

⁵The preprocessing codes of EMERGE [53] and KerPrint [50] are not public, and the reported results for these models use incorrect preprocessing, so we do not include their experimental results.

5.2 Main Results

Table 1 compares the performance of mortality prediction and readmission prediction for TRACER and the latest baselines on MIMIC-III. TRACER outperforms all competitors across all metrics on both tasks. For the mortality prediction task, sensitivity and Macro F1 score improve by 28.5% and 19.7%, respectively, compared to the state-of-the-art method, i.e., KARE. We speculate that this gap may be due to the patient-specific trajectories sampled solely from our severity-weighted medical KG. Table 2 shows the performance of the compared methods on MIMIC-IV. For the mortality prediction task, TRACER achieves the state of the art performance across all metrics except specificity, though it trails the best competitor by only 0.8%. In contrast, for the readmission prediction task, TRACER surpasses the competitors on specificity by at least 15.5%. These results demonstrate that TRACER is highly competitive for both tasks. In addition, the observed performance gains of TRACER are consistent across different model categories. Specifically, models relying solely on EHR sequences (e.g., GRASP, StageNet, RAM-EHR) are consistently outperformed by approaches that incorporate external knowledge or additional features, such as medical KGs (e.g., KARE) or retrieved medical concepts (e.g., RAM-EHR, MedRetriever). TRACER further advances this trend by leveraging severity-weighted KG trajectories and clinical text retrieval, resulting in a clear and monotonic improvement from these existing methods. Importantly, the improvement achieved by TRACER substantially exceeds the typical performance gains observed in prior work. Across existing baselines, the average incremental improvement in Macro F1 over preceding methods is approximately 2.0 points, whereas TRACER improves Macro F1 by 18.2 points over the strongest prior model (KARE). This indicates that TRACER delivers a substantive methodological advance beyond incremental model refinements. All reported results are averaged over three runs; standard deviations are below 0.06 across the main variants, indicating stable trends. More detailed analysis on mortality and readmission prediction performances are included in the extended version [14].

5.3 Ablation Study

We conduct an ablation study to examine the contribution of each component in TRACER. To clarify attribution in the multi-component pipeline, we organize the variants around core modules: SMKG, trajectory retrieval, clinical note retrieval, and similar-patient retrieval. We also evaluate the variants with combined modules removed to assess whether the overall performance is driven by a single component or by complementary integration. All ablation results are averaged over three runs, and the standard deviation of each variant is below 0.06, indicating stable trends across runs.

The variant “w/o Sev.” samples trajectories from the medical KG without severity scores; “w/o Risk.”, “w/o Prot.”, and “w/o Traj.” denote the variants in which the LLM performs reasoning based on all the proposed supporting evidence except risk trajectories, except protective trajectories, and except both trajectories, respectively; “w/o Pat.” denotes the variant that builds the proposed supporting evidence except only the trajectories and demographics of the top- m similar patients, and “w/o Clin.” represents the variant with all the supporting evidence except top- n clinical notes. We also evaluate a

variant with the patient medical profile (w/o PMP) removed or that with both PMP and top- m patients removed (w/o (PMP+Pat.)).

Since MIMIC-III and MIMIC-IV are highly imbalanced, sensitivity is the most important metric [16] because failing to identify true positive cases such as actual mortality or readmission can have critical clinical consequences.

Table 3 shows that the three core modules contribute complementary gains. Removing severity from SMKG reduces mortality sensitivity from 0.7128 to 0.4950, confirming the role of severity grounding. Trajectory retrieval is the most influential component for mortality prediction, as removing all trajectories reduces Macro F1 from 0.8197 to 0.6302. In particular, removing risk trajectories causes the largest sensitivity drop, indicating that severity-aware disease progression is critical for detecting mortality cases. In contrast, readmission prediction is more sensitive to clinical notes and similar-patient evidence. Removing clinical notes decreases Macro F1 from 0.8709 to 0.7981 and sensitivity from 0.7809 to 0.6762, suggesting that unstructured notes capture short-term deterioration and discharge-related cues that are not fully represented by structured codes. Removing similar patients causes a clear performance drop, indicating that cohort-level context provides useful complementary evidence. Finally, the combined ablations show that TRACER is not driven by a single component, but by the coordinated use of SMKG, trajectory retrieval, clinical notes, and similar-patient evidence.

6 Analysis

6.1 Consistency of Severity Score Generation

To clarify the consistency of severity score generation, we generate each disease-level severity score three times using different LLM backbones and evaluate the resulting end-to-end pipeline on MIMIC-III for both mortality and readmission prediction. For each disease, we also calculate the standard deviation, dubbed Severity S.D., from the mean of the three generated scores. Lower values indicate more consistent severity estimation across LLM backbones.

Table 4 summarizes end-to-end prediction performance and severity-score consistency across different LLM backbones. Severity S.D. decreases from 2.02 with GPT-4o-mini to 1.59 with GPT-5-nano and 1.16 with GPT-o3, indicating that larger models generate more consistent severity scores for the same disease. A similar trend is observed in both mortality and readmission prediction performance, where larger models achieve higher accuracy. In particular, GPT-o3 more effectively exploits severity-aware trajectories and retrieved clinical notes, resulting in over a 3% improvement in Macro F1 for mortality prediction compared to GPT-4o-mini. Importantly, severity variability remains low even for GPT-4o-mini, demonstrating that the proposed severity scoring procedure is robust and not overly sensitive to the choice of LLM.

6.2 Effectiveness of SMKG

To evaluate the effectiveness of SMKG, we replace the UMLS-based medical knowledge graph used in MedGraphRAG [43] with the proposed SMKG, and analyze the resulting performance changes. Experiments are conducted on three representative medical question answering benchmarks: MedQA [17], MedMCQA [36], and PubMedQA [19]. Table 5 show that substituting UMLS with SMKG

Table 4: Performance comparison on MIMIC-III across different LLMs for mortality and readmission prediction, along with severity score variability, where lower values indicate more consistent severity score generation.

LLM	Mortality Prediction				Readmission Prediction				Severity S.D.
	Accuracy	Macro F1	Sensitivity	Specificity	Accuracy	Macro F1	Sensitivity	Specificity	
GPT-4o-mini	0.9275	0.8197	0.7128	0.9531	0.9135	0.8709	0.7809	0.9501	2.02
GPT-5-nano	0.9356	0.8289	0.7311	0.9608	0.9222	0.8805	0.7963	0.9598	1.59
GPT-o3	0.9460	0.8447	0.7562	0.9701	0.9471	0.8916	0.8186	0.9701	1.16

Table 5: Performance comparison of MedGraphRAG when replacing the UMLS graph with the proposed SMKG across multiple medical QA benchmarks.

Methods	MedMCQA	PubMedQA	MedQA
MedGraphRAG - UMLS	81.5	83.3	91.3
MedGraphRAG - SMKG	84.2	86.5	94.9

Table 6: Comparison of inference efficiency between TRACER and KARE on the MIMIC-III for mortality prediction task.


Model	Inference Time (s)	Patient-specific Context Retrieval Time (s)	Macro F1	Sensitivity
KARE	117	92	0.6375	0.3711
TRACER	88	60	0.8284	0.6831

consistently improves performance across all benchmarks. In particular, on the MedQA dataset, the accuracy increases by 3.94 percentage points, representing the most substantial gain among the evaluated datasets. These findings indicate that, unlike conventional medical knowledge graphs that treat all disease concepts with equal importance, assigning severity scores to disease nodes enables the model to more effectively distinguish relative clinical risk and importance among diseases. Overall, SMKG enhances medical question answering performance by explicitly encoding disease severity into the graph structure, allowing the model to prioritize clinically higher-risk conditions and retrieve more relevant evidence during reasoning.

6.3 Inference efficiency

Table 6 compares the inference time and accuracy of TRACER and KARE, i.e., the best competitor in terms of overall prediction performance⁶. The results show that TRACER is approximately 33% faster than KARE in terms of inference time, while achieving 29.9% and 84.1% higher Macro F1 and sensitivity, respectively. Both methods perform patient-specific context retrieval to support reasoning-based prediction, but differ substantially in how such information is retrieved. KARE performs patient subgraph retrieval by iteratively expanding multi-hop neighborhoods from patient-related medical concepts in its medical knowledge graph, while TRACER retrieves severity-aware, temporally ordered disease trajectories

⁶For TRACER, inference time is measured as the total duration of the inference stage. For KARE, inference time is defined as the sum of the durations required for positive/negative patient retrieval, reasoning chain construction, and LLM prediction.

**Figure 5: Sensitivity analysis of the number of similar patients, the number of clinical notes, the proportion of trajectories, and the weight for the MMR score on MIMIC-III.**

tailored to the target patient from the SMKG. As shown in Table 6, the patient subgraph retrieval time for KARE is 92 seconds, while the patient trajectory retrieval for TRACER requires only 60 seconds, making TRACER about 35% faster. This demonstrates that leveraging trajectories offers clear advantages in both efficiency and predictive accuracy.

6.4 Hyperparameter Sensitivity

We evaluate the impact of hyperparameters on mortality prediction performance on MIMIC-III. Figure 5 shows the changes in Macro F1 score and sensitivity for: (i) the number m of similar patients, (ii) the number n of clinical notes, (iii) the proportion ρ of trajectories, and (iv) weight factor λ_{MMR} for MMR score. As described in Figure 5, m and n exhibit similar trends, achieving the highest performance when both are set to 5, with Macro F1 score and sensitivity reaching 0.8284 and 0.6831, respectively. When set below 5, the performance decreases slightly, while values above 5 lead to a drop in performance due to excessive information. The weighting factor λ_{MMR} for the MMR score achieves the best performance at 0.7, but shows a decline when set above 0.7. In addition, the proportion ρ of trajectories most relevant to each patient and each (risk or protective) factor yields the best results at 5, while performance slightly drops for values less than 5. Further details are presented in the extended version [14].

Table 7: Performance of different methods selecting similar patients on MIMIC-IV. The best results are denoted in bold.

Methods	Mortality		Readmission	
	Macro F1	Sensitivity	Macro F1	Sensitivity
StageNet embedding	0.8798	0.7696	0.7907	0.8732
Last-visit Jaccard	0.8787	0.7646	0.7904	0.8712
Recent-visit Jaccard	0.8830	0.7722	0.7942	0.8752

Table 8: Robustness analysis across different LLM backbones on MIMIC-III. The best results are denoted in bold.

Methods	Macro F1	Sensitivity	Specificity	Accuracy
GPT-4o-mini	0.8197	0.7128	0.9531	0.9275
LLaMA2-8B	0.7729	0.5446	0.9638	0.9182
Qwen3.5-2B	0.7581	0.5842	0.9434	0.9045

6.5 Similar Patient Retrieval Strategy

We conduct experiments to compare the performance of mortality and readmission prediction according to the criterion used to retrieve similar patients. As shown in Table 7, we compare three strategies: (i) StageNet embedding, which retrieves similar patients using cosine similarity over the entire visit-sequence embedding; (ii) Last-visit Jaccard, which uses Jaccard similarity over medical concepts appearing only in the last historical visit; and (iii) Recent-visit Jaccard, which uses Jaccard similarity over medical concepts appearing in the most recent two historical visits.

The results show that Recent-visit Jaccard achieves the best performance across both tasks. This supports our design choice of retrieving similar patients based on exact recent overlap rather than broad sequence-level semantic similarity. In clinical risk prediction, semantically related concepts may still imply different prognostic risks, e.g., sepsis and septic shock are closely related but differ substantially in mortality risk. Beyond predictive performance, Recent-visit Jaccard achieves an average retrieval speed per patient that is approximately 19% faster than the StageNet baseline. Balancing both effectiveness and efficiency, Recent-visit Jaccard emerges as the most practical choice⁷.

6.6 Robustness to LLM Backbones

Table 8 evaluates the robustness of TRACER across different LLM backbones on MIMIC-III mortality prediction. GPT-4o-mini achieves the best overall performance, while the open-source LLMs also maintain competitive results. In particular, LLaMA2-8B obtains the highest specificity, and Qwen3.5-2B preserves reasonable sensitivity and accuracy despite its smaller scale. These results indicate that TRACER is not specific to GPT-family models, but can remain effective when instantiated with different LLM backbones. Additional robustness analysis across knowledge snapshots is provided in Appendix I.2 of the extended version [14].

⁷Similar-patient evidence is used only as auxiliary context: the final prediction remains primarily grounded in the target patient’s own trajectories, clinical notes, and demographics, thereby reducing the risk of shortcut reasoning based only on peer outcomes.

Table 9: Hop-limit analyses on MIMIC-III mortality prediction.

Analysis	Method	Macro-F1	Sensitivity	Specificity	Accuracy
Hop limit	1-hop	0.7960	0.6732	0.9471	0.9177
	2-hop	0.8083	0.6930	0.9506	0.9229
	3-hop	0.8197	0.7128	0.9531	0.9275
	4-hop	0.8194	0.7029	0.9542	0.9291

6.7 Hop-limit Analyses

We evaluate the impact of the maximum path length h during trajectory extraction to determine the optimal amount of KG context. As shown in Table 9, a restrictive hop limit risks omitting clinically relevant paths, whereas a large hop limit can introduce weakly related or noisy connections. The results demonstrate that predictive performance improves from $h = 1$ to $h = 3$, with Macro F1 rising from 0.7960 to 0.8197 and sensitivity increasing from 0.6732 to 0.7128. Extending the search to $h = 4$ hops results in a slight degradation in sensitivity, confirming that over-expansion introduces distracting or irrelevant evidence to the LLM. Consequently, setting $h = 3$ strikes the most effective practical balance between evidence coverage and noise mitigation in trajectory retrieval.

7 Conclusion

In this paper, we propose **TRACER**, a severity-aware trajectory RAG framework designed to enhance mortality and readmission prediction. By harnessing the reasoning capabilities of LLMs, TRACER enriches a patient’s visit sequence with highly relevant medical context. Central to this framework is our severity-weighted medical KG, which assigns literature-backed severity scores to diagnosis concepts, thereby emphasizing high-risk clinical signals. We demonstrate the effectiveness of augmenting a patient’s history with three critical dimensions of the patient: (1) severity-aware clinical trajectories that can capture disease progression across visits, (2) similar peer cases based on recent visit patterns, and (3) contextual evidence from unstructured clinical notes. This integration significantly improves both predictive performance and clinical transparency. Extensive experiments on MIMIC-III and MIMIC-IV demonstrate that TRACER consistently outperforms state-of-the-art baselines in sensitivity, Macro F1, specificity, and accuracy, offering a robust solution for interpretable, data-driven clinical risk assessment.

Acknowledgments

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2020-II201373, Artificial Intelligence Graduate School Program(Hanyang University)) and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2026-25487321), to K. Jeon, Y. Ko, and H. Kim. The work of W. Jung was supported by the National R&D Program for Cancer Control through the National Cancer Center(NCC) funded by the Ministry of Health & Welfare, Republic of Korea (RS-2025-02264000).

References

- [1] Ahmed Allam, Mate Nagy, George Thoma, and Michael Krauthammer. 2019. Neural networks versus Logistic regression for 30 days all-cause readmission prediction. *Sci. Rep.* 9, 1 (June 2019), 9277.
- [2] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. arXiv:1803.01271 [cs.LG] <https://arxiv.org/abs/1803.01271>
- [3] Olivier Bodenreider. 2004. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research* 32, suppl_1 (2004), D267–D270.
- [4] Kathi Canese and Sarah Weis. 2013. *PubMed: the bibliographic database* (2nd ed.). National Center for Biotechnology Information (US), Bethesda, MD, Chapter 2(1). <https://www.ncbi.nlm.nih.gov/books/NBK153385/> Created Oct 9 2002; Updated Mar 20 2013.
- [5] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Alessandro Moschitti, Bo Pang, and Walter Daelemans (Eds.). Association for Computational Linguistics, Doha, Qatar, 1724–1734. doi:10.3115/v1/D14-1179
- [6] Edward Choi, Mohammad Taha Bahadori, Joshua A. Kulas, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. 2016. RETAIN: an interpretable predictive model for healthcare using reverse time attention mechanism. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 3512–3520.
- [7] Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, and Jimeng Sun. 2017. GRAM: Graph-based Attention Model for Healthcare Representation Learning. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Halifax, NS, Canada) (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 787–795. doi:10.1145/3097983.3098126
- [8] P. Chung, C. T. Fong, A. M. Walters, N. Aghaeepour, M. Yetisgen, and V. N. O'Reilly-Shah. 2024. Large Language Model Capabilities in Perioperative Risk Prediction and Prognostication. *JAMA Surgery* 159, 8 (2024), 928–937. doi:10.1001/jamasurg.2024.1621
- [9] Sajad Darabi, Mohammad Kachuee, Shayan Fazeli, and Majid Sarrafzadeh. 2020. TAPER: Time-Aware Patient EHR Representation. *IEEE Journal of Biomedical and Health Informatics* 24, 11 (2020), 3268–3275. doi:10.1109/JBHI.2020.2984931
- [10] Junyi Gao, Cao Xiao, Yasha Wang, Wen Tang, Lucas M. Glass, and Jimeng Sun. 2020. StageNet: Stage-Aware Neural Networks for Health Risk Prediction. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 530–540. doi:10.1145/3366423.3380136
- [11] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6894–6910. doi:10.18653/v1/2021.emnlp-main.552
- [12] Kexin Huang, Jaan Altsaar, and Rajesh Ranganath. 2020. ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission. arXiv:1904.05342 [cs.CL] <https://arxiv.org/abs/1904.05342>
- [13] W. Huang, A. Nurhafizah, A. Frederich, A. R. Khairunnisa, C. Kezia, M. I. Fathoni, S. Samban, and S. Flindy. 2025. Risk and Protective Factors of Poor Clinical Outcomes in Heart Failure with Improved Ejection Fraction Population: A Systematic Review and Meta-Analysis. *Current Cardiology Reports* 27, 1 (2025), 4. doi:10.1007/s11886-024-02180-w
- [14] Kyunghoon Jeon, Yoomin Ko, Woohwan Jung, and Hyunjoon Kim. 2026. TRACER: Extended Version. <https://github.com/KyunghoonJeon/TRACER.git>. Accessed: 2026-06-02.
- [15] Pengcheng Jiang, Cao Xiao, Adam Richard Cross, and Jimeng Sun. 2024. GraphCare: Enhancing Healthcare Predictions with Personalized Knowledge Graphs. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=tVTN7Zs0ml>
- [16] Pengcheng Jiang, Cao Xiao, Minhao Jiang, Parminder Bhatia, Taha Kass-Hout, Jimeng Sun, and Jiawei Han. 2025. Reasoning-Enhanced Healthcare Predictions with Knowledge Graph Community Retrieval. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=8fLgt7PQza>
- [17] Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams. *Applied Sciences* 11, 14 (2021). doi:10.3390/app11146421
- [18] Jiajie Jin, Yutao Zhu, Yujia Zhou, and Zhicheng Dou. 2024. BIDER: Bridging Knowledge Inconsistency for Efficient Retrieval-Augmented LLMs via Key Supporting Evidence. In *Findings of the Association for Computational Linguistics: ACL 2024*. arXiv:2402.12174 [cs.CL] doi:10.48550/arXiv.2402.12174
- [19] Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. PubMedQA: A Dataset for Biomedical Research Question Answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2567–2577. doi:10.18653/v1/D19-1259
- [20] Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data* 10, 1 (2023), 1.
- [21] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9.
- [22] Helena Chmura Kraemer, Alan E. Kazdin, David R. Offord, Ronald C. Kessler, Peter S. Jensen, and David J. Kupfer. 1997. Coming to Terms With the Terms of Risk. *Archives of General Psychiatry* 54, 4 (1997), 337–343. doi:10.1001/archpsyc.1997.01830160065009
- [23] G. A. Lazarin, F. Hawthorne, N. S. Collins, E. A. Platt, E. A. Evans, and I. S. Haque. 2014. Systematic Classification of Disease Severity for Evaluation of Expanded Carrier Screening Panels. *PLoS ONE* 9, 12 (2014), e114391. doi:10.1371/journal.pone.0114391
- [24] Yue Ling. 2023. Bio+Clinical BERT, BERT Base, and CNN Performance Comparison for Predicting Drug-Review Satisfaction. ArXiv abs/2308.03782 (2023). <https://api.semanticscholar.org/CorpusID:260704314>
- [25] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173. doi:10.1162/tacl_a_00638
- [26] Chang Lu, Tian Han, and Yue Ning. 2022. Context-aware health event prediction via transition functions on dynamic disease graphs. *Proc. Conf. AAAI Artif. Intell.* 36, 4 (June 2022), 4567–4574.
- [27] Junyu Luo, Muchao Ye, Cao Xiao, and Fenglong Ma. 2020. HiTANet: Hierarchical Time-Aware Attention Networks for Risk Prediction on Electronic Health Records. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Virtual Event, CA, USA) (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 647–656. doi:10.1145/3394486.3403107
- [28] Weimin Lyu, Xinyu Dong, Rachel Wong, Songzhu Zheng, Kayley Abell-Hart, Fusheng Wang, and Chao Chen. 2022. A multimodal transformer: Fusing clinical notes with structured EHR data for interpretable in-hospital mortality prediction. *AMIA Annu. Symp. Proc.* 2022 (2022), 719–728.
- [29] Fenglong Ma, Quanzeng You, Houping Xiao, Radha Chitta, Jing Zhou, and Jing Gao. 2018. KAME: Knowledge-based Attention Model for Diagnosis Prediction in Healthcare. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (Torino, Italy) (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 743–752. doi:10.1145/3269206.3271701
- [30] Liantao Ma, Junyi Gao, Yasha Wang, Chaohe Zhang, Jiangtao Wang, Wenjie Ruan, Wen Tang, Xin Gao, and Xinyu Ma. 2020. AdaCare: Explainable Clinical Health Status Representation Learning via Scale-Adaptive Feature Extraction and Recalibration. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (Apr. 2020), 825–832. doi:10.1609/aaai.v34i01.5427
- [31] Liantao Ma, Chaohe Zhang, Yasha Wang, Wenjie Ruan, Jiangtao Wang, Wen Tang, Xinyu Ma, Xin Gao, and Junyi Gao. 2020. ConCare: Personalized Clinical Feature Embedding via Capturing the Healthcare Context. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (Apr. 2020), 833–840. doi:10.1609/aaai.v34i01.5428
- [32] Parvati Naliyathaliyazhachayil, Raajitha Muthyala, Judy Wawira Gichoya, and Saptarshi Purkayastha. 2025. Evaluating the Reasoning Capabilities of Large Language Models for Medical Coding and Hospital Readmission Risk Stratification: Zero-Shot Prompting Approach. In *Journal of Medical Internet Research*, Vol. 27. e74142. doi:10.2196/74142
- [33] P. Naliyathaliyazhachayil, R. Muthyala, J. W. Gichoya, and S. Purkayastha. 2025. Evaluating the Reasoning Capabilities of Large Language Models for Medical Coding and Hospital Readmission Risk Stratification: Zero-Shot Prompting Approach. *Journal of Medical Internet Research* 27 (2025), e74142. doi:10.2196/74142
- [34] Phuoc Nguyen, Truyen Tran, Nilmini Wickramasinghe, and Svetha Venkatesh. 2017. DeepIR: A Convolutional Net for Medical Records. *IEEE Journal of Biomedical and Health Informatics* 21, 1 (2017), 22–30. doi:10.1109/JBHI.2016.2633963
- [35] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, et al. 2024. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL] <https://arxiv.org/abs/2303.08774>
- [36] Ankit Pal, Logesh Kumar Umashathi, and Malaikannan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*. PMLR, 248–260.
- [37] Chantal Pellegrini, Ege Özsoy, David Bani-Harouni, Matthias Keicher, and Nassir Navab. 2025. From EHRs to Patient Pathways: Scalable Modeling of Longitudinal Health Trajectories with LLMs. arXiv:2506.04831 [cs.LG] <https://arxiv.org/abs/>

- 2506.04831
- [38] Ofir Ben Shoham and Nadav Rappoport. 2024. CPLLM: Clinical Prediction with Large Language Models. arXiv:2309.11295 [cs.CL]. <https://arxiv.org/abs/2309.11295>
 - [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 6000–6010.
 - [40] Xiaohua Wang, Zhenghua Wang, Xuan Gao, Feiran Zhang, Yixin Wu, Zhibo Xu, Tianyuan Shi, Zhengyuan Wang, Shizheng Li, Qi Qian, Ruicheng Yin, Changze Lv, Xiaoqing Zheng, and Xuanjing Huang. 2024. Searching for Best Practices in Retrieval-Augmented Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 17716–17736. doi:10.18653/v1/2024.emnlp-main.981
 - [41] Zixiang Wang, Yinghao Zhu, Junyi Gao, Xiaochen Zheng, Yuhui Zeng, Wen Tang, Yasha Wang, Ewen M Harrison, Chengwei Pan, Liantao Ma, and Ling Wang. 2024. RetCare: Towards Interpretable Clinical Decision Making through LLM-Driven Medical Knowledge Retrieving. In *Artificial Intelligence and Data Science for Healthcare: Bridging Data-Centric AI and People-Centric Healthcare*. <https://openreview.net/forum?id=jqo1vk63qE>
 - [42] Zixiang Wang, Yinghao Zhu, Huiya Zhao, Xiaochen Zheng, Dehao Sui, Tianlong Wang, Wen Tang, Yasha Wang, Ewen Harrison, Chengwei Pan, Junyi Gao, and Liantao Ma. 2025. ColaCare: Enhancing Electronic Health Record Modeling through Large Language Model-Driven Multi-Agent Collaboration. In *Proceedings of the ACM on Web Conference 2025* (Sydney NSW, Australia) (WWW '25). Association for Computing Machinery, New York, NY, USA, 2250–2261. doi:10.1145/3696410.3714877
 - [43] Junde Wu, Jiayuan Zhu, Yunli Qi, Jingkun Chen, Min Xu, Filippo Menolascina, Yueming Jin, and Vicente Grau. 2025. Medical Graph RAG: Evidence-based Medical Large Language Model via Graph Retrieval-Augmented Generation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 28443–28467. doi:10.18653/v1/2025.acl-long.1381
 - [44] Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. Benchmarking Retrieval-Augmented Generation for Medicine. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 6233–6251. doi:10.18653/v1/2024.findings-acl.372
 - [45] Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. RECOMP: Improving Retrieval-Augmented LMs with Compression and Selective Augmentation. arXiv:2310.04408 [cs.CL]. <https://arxiv.org/abs/2310.04408>
 - [46] Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Bowen Jin, May Dongmei Wang, Joyce Ho, and Carl Yang. 2024. RAM-EHR: Retrieval Augmentation Meets Clinical Predictions on Electronic Health Records. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 754–765. doi:10.18653/v1/2024.acl-short.68
 - [47] Yongxin Xu, Xu Chu, Kai Yang, Zhiyuan Wang, Peinie Zou, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. 2023. SeqCare: Sequential Training with External Medical Knowledge Graph for Diagnosis Prediction in Healthcare Data. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 2819–2830. doi:10.1145/3543507.3583543
 - [48] Yongxin Xu, Xinke Jiang, Xu Chu, Rihong Qiu, Yujie Feng, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. 2025. DearLLM: Enhancing Personalized Healthcare via Large Language Models-Deduced Feature Correlations. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 1 (Apr. 2025), 941–949. doi:10.1609/aaai.v39i1.32079
 - [49] Fan Yang, Jian Zhang, Wanyi Chen, Yongxuan Lai, Ying Wang, and Quan Zou. 2022. DeepMPM: a mortality risk prediction model using longitudinal EHR data. *BMC Bioinformatics* 23, 1 (Oct. 2022), 423.
 - [50] Kai Yang, Yongxin Xu, Peinie Zou, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. 2023. KerPrint: Local-Global Knowledge Graph Enhanced Diagnosis Prediction for Retrospective and Prospective Interpretations. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 4 (Jun. 2023), 5357–5365. doi:10.1609/aaai.v37i4.25667
 - [51] Muchao Ye, Suhan Cui, Yaqing Wang, Junyu Luo, Cao Xiao, and Fenglong Ma. 2021. MedRetriever: Target-Driven Interpretable Health Risk Prediction via Retrieving Unstructured Medical Text. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (Virtual Event, Queensland, Australia) (CIKM '21). Association for Computing Machinery, New York, NY, USA, 2414–2423. doi:10.1145/3459637.3482273
 - [52] Chaohe Zhang, Xin Gao, Liantao Ma, Yasha Wang, Jiangtao Wang, and Wen Tang. 2021. GRASP: Generic Framework for Health Status Representation Learning

Based on Incorporating Knowledge from Similar Patients. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 1 (May 2021), 715–723. doi:10.1609/aaai.v35i1.16152

- [53] Yinghao Zhu, Changyu Ren, Zixiang Wang, Xiaochen Zheng, Shiyun Xie, Junlan Feng, Xi Zhu, Zhoujun Li, Liantao Ma, and Chengwei Pan. 2024. EMERGE: Enhancing Multimodal Electronic Health Records Predictive Modeling with Retrieval-Augmented Generation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management* (Boise, ID, USA) (CIKM '24). Association for Computing Machinery, New York, NY, USA, 3549–3559. doi:10.1145/3627673.3679582
- [54] Yinghao Zhu, Zixiang Wang, Long He, Shiyun Xie, Xiaochen Zheng, Liantao Ma, and Chengwei Pan. 2024. PRISM: Mitigating EHR Data Sparsity via Learning from Missing Feature Calibrated Prototype Patient Representations. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management* (Boise, ID, USA) (CIKM '24). Association for Computing Machinery, New York, NY, USA, 3560–3569. doi:10.1145/3627673.3679521

A Robustness to Data Sparsity across Visit Sequence Lengths

We analyze the robustness under increasing data sparsity that arises as visit sequences become longer. Patients are divided into five groups according to the exact number of visits: 1, 2, 3, 4, and 5. Figure 6 illustrates the performance trends across visit lengths, where the number of visits is used as the horizontal axis, the line plots present the Macro F1 and sensitivity of the baselines and TRACER, and the background bars represent the percentage of patients associated with each visit length.

Most baselines exhibit a steep decline in Macro F1 as the number of visits increases, indicating that their performance degrades rapidly for patients with longer visit sequences. In contrast, TRACER shows a much more gradual decrease in Macro F1, maintaining comparatively stable performance even in the longer-visit groups. A similar trend is observed for sensitivity: as the number of visits increases, TRACER demonstrates a smooth and consistent increase, whereas baselines exhibit more abrupt or unstable changes.

These results show that TRACER is more robust to the data sparsity that arises with longer visit sequences. By leveraging trajectories of similar patients selected based on the target patient's recent visits, TRACER focuses on the most relevant and up-to-date clinical status rather than simply accumulating historical records. This design enables more stable and reliable predictions for patients with longer visit histories.

B Impact of Repacking and Refinement

Table 10 validates the effectiveness of trajectory refinement and repacking. In the variant "w/o Traj. Ref.", all shortest paths with up to three hops between concepts in adjacent visits form the trajectories. In the "w/o Traj. Rep." variant, all trajectories are randomly shuffled prior to input to the LLM. In contrast, trajectory refinement selects only the most important paths based on the MMR score. Experimental results show that the removal of trajectory refinement decreases the sensitivity of mortality prediction by 7.2% (0.6831 → 0.6336).

This demonstrates that MMR-based refinement has a significant positive effect on performance. Furthermore, omitting trajectory repacking results in a 5.8% decrease in sensitivity (0.6831 → 0.6435), while the removal of similar patient repacking leads to a 4.3% decrease (0.6831 → 0.6534). This demonstrates that repacking trajectories is more effective than repacking similar patients, indicating that the LLM is more sensitive to the order and arrangement of

Table 10: Ablation study on reranking and repacking in TRACER


Models	Mortality		Readmission	
	Macro F1	Sensitivity	Macro F1	Sensitivity
TRACER	0.8284	0.6831	0.8709	0.7809
w/o Traj. Ref.	0.7944	0.6336	0.8469	0.7476
w/o Traj. Rep.	0.8097	0.6435	0.8626	0.7714
w/o Pat. Rep.	0.8153	0.6534	0.8604	0.7571

Table 11: Effect of severity score granularity on mortality prediction performance on the MIMIC-III dataset.

Severity Scale	Macro F1	Sensitivity	Specificity	Accuracy
1–100 (5-level prompt)	0.8284	0.6831	0.9639	0.9334
1–20	0.8197	0.7128	0.9531	0.9275
1–5	0.7271	0.5842	0.9194	0.8831

Table 12: Coverage-matched comparison with GraphCare on MIMIC-III mortality prediction. Coverage denotes the amount of retrieved KG evidence exposed to the predictor.

Method	Macro F1	Sensitivity	Specificity	Accuracy
TRACER	0.8024	0.6831	0.9482	0.9206
GraphCare	0.5963	0.2552	0.9577	0.9229

**Figure 6: Performance in terms of Macro F1 and sensitivity across visit sequence lengths for the readmission prediction task, with the percentage of patients by number of visits.**

target patient’s trajectories. Thus, careful selection and ordering of trajectories are essential for optimal model performance.

C Impact of Severity Score Granularity.

We further investigate the impact of severity score granularity on mortality prediction performance. Specifically, we compare three scoring schemes: a fine-grained severity score in the range of 1–100 (implemented as a 5-level prompt), a moderately coarse scheme in the range of 1–20, and a coarse-grained scheme in the range of 1–5.

As shown in Table 11, the 1–100 severity score achieves the best performance in terms of Macro F1, specificity, and accuracy, while the 1–20 scheme yields the highest sensitivity. In contrast, the 1–5 scheme consistently underperforms across all metrics, with

a particularly notable drop in Macro F1 and sensitivity. This performance degradation suggests that coarse-grained scoring fails to capture meaningful differences among patients with varying levels of clinical severity. In additional experiments, we observe that severity scores are more densely distributed in the 51–100 range than in the 1–50 range, indicating that fine-grained scores better differentiate high-risk patients. This observation aligns with the design motivation of clinical risk stratification and is consistent with prior severity indices, such as the Charlson Comorbidity Index (CCI), which also emphasizes higher-weight severe conditions. These results indicate that fine-grained severity modeling is crucial for accurately capturing disease severity and improving mortality prediction, and that overly coarse scoring schemes may obscure clinically important risk differences.

D Correlation Analysis between Severity Scores and Mortality Labels.

For in-depth analysis, we compute the point-biserial correlation coefficient between a patient-level severity score and mortality label for every patient. Since our severity scores are defined at the diagnosis level, we aggregate them into the patient-level severity score by averaging the severity scores of all diagnoses observed in the patient’s visit history, similar to the scoring technique of Charlson Comorbidity Index (CCI; normalized to a 37-point scale). The patient-level severity scores (normalized to a 100-point scale) achieve a correlation of 0.38 with the mortality labels, while the CCI shows a correlation of 0.29. Both correlations are statistically significant with p-values below 0.05. We observe that CCI tends to assign higher scores to patients with specific high-weight comorbid conditions, as it relies on predefined disease-specific weights. In contrast, the proposed severity score captures a broader range of clinical severity signals, which leads to a stronger overall association with mortality outcomes. Notably, the mortality label in our dataset is highly imbalanced, with only 10.83% positive cases. This label skew generally suppresses the magnitude of correlation coefficients, suggesting that the observed correlations may underestimate the true relationship between clinical severity and mortality risk.

E Comparison under the Same Amount of Retrieved KG Evidence

We further compare TRACER with GraphCare while giving both methods a similar amount of retrieved KG evidence. This experiment tests whether TRACER performs better simply because it retrieves more KG information. Specifically, we control the number of KG triples or paths provided to the prediction module so that the two methods receive a comparable amount of evidence. As shown in Table 12, TRACER still achieves higher Macro F1 and sensitivity than GraphCare. This result suggests that TRACER’s improvement comes from selecting more patient-relevant trajectory evidence, rather than from using a larger amount of KG evidence.