

# Knowledge Graph Augmented Large Language Models for Disease Prediction

Ruiyu Wang<sup>1</sup>, Tuan Vinh<sup>2</sup>, Ran Xu<sup>1</sup>, Yuyin Zhou, PhD<sup>3</sup>, Jiaying Lu, PhD<sup>4</sup>, Francisco Pasquel, MD<sup>5</sup>, Mohammed K Ali, MD<sup>6</sup>, Carl Yang, PhD<sup>1,4,6</sup>

<sup>1</sup> Department of Computer Science, Emory University, Atlanta, GA, USA

<sup>2</sup> Division of Medical Sciences, Oxford University, Oxford, UK

<sup>3</sup> Department of Computer Science and Engineering, UCSC, CA, USA


<sup>4</sup> Nell Hodgson Woodruff School of Nursing, Emory University, Atlanta, GA, USA

<sup>5</sup> School of Medicine, Emory University, Atlanta, GA, USA

<sup>6</sup> Rollins School of Public Health, Emory University, Atlanta, GA, USA

## Abstract

Electronic health records (EHRs) support strong clinical prediction but often provide coarse, post hoc explanations that are hard to use for patient-level decisions. We propose a knowledge-graph (KG)-guided chain-of-thought (CoT) framework for visit-level disease prediction on MIMIC-III. We map ICD-9 codes to PrimeKG, mine disease-relevant nodes and paths, and use these paths to scaffold temporally consistent CoT explanations, retaining only samples whose conclusions match observed outcomes. We then fine-tune lightweight LLaMA-3.1-Instruct-8B and Gemma-7B models on two small cohorts (400 and 1,000 index visits) across ten PrimeKG-mapped diseases. Our models outperform strong classical baselines, reaching AUROC of 0.66–0.70 and macro-AUPR of 0.40–0.47. Without additional training, the models transfer zero-shot to the CRADLE cohort, improving accuracy from 0.40–0.51 to 0.72–0.77. Blinded clinicians consistently prefer KG-guided CoT for clarity, relevance, and correctness. Code is available at <https://github.com/JonathanWry/KG-guided-LLM-pipeline>.



## Introduction

Electronic health records (EHRs) are now routinely collected across large and diverse healthcare systems, capturing longitudinal data on diagnoses, medications, and laboratory results from millions of patients.<sup>1,2,3</sup> As interest in digital medicine grows,<sup>4,5</sup> these data are increasingly used to train predictive models for outcomes such as in-hospital mortality, readmission, and disease onset, often using benchmark cohorts like MIMIC-III.<sup>6,7</sup> Classical models such as logistic regression and tree-based ensembles remain popular due to their robustness, global interpretability, and ease of deployment in clinical systems,<sup>7</sup> while neural architectures including recurrent and attention-based models have been explored to capture temporal dynamics and co-occurrence patterns in EHR data.<sup>8,2</sup> Despite these advances in predictive performance, most EHR models provide only coarse, post hoc interpretability (such as global feature importance, attribution scores, or attention weights), offering limited support for patient-specific, clinician-facing reasoning about why a particular outcome is likely.<sup>7,8,9,10</sup>

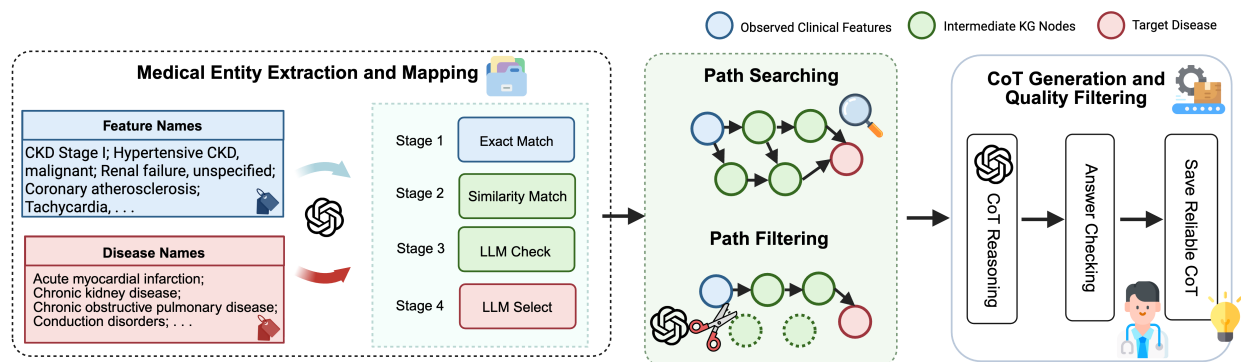
Advances in large language models (LLMs) and reasoning-oriented architectures offer a potential remedy: LLMs can generate intermediate reasoning steps alongside predictions, enabling explicit chain-of-thought (CoT) narratives that articulate why a given prediction should hold.<sup>11,12,13,14</sup> In principle, such step-by-step explanations could bridge the gap between accurate prediction and clinically usable justification for clinical predictions. However, applying CoT in medicine remains challenging. High-quality, clinically validated CoT datasets are scarce<sup>15,16</sup> and unconstrained LLM reasoning can introduce hallucinated or medically incorrect intermediate steps, which is particularly problematic in high-stakes clinical decision making.<sup>17,18</sup> These limitations have motivated efforts to constrain or ground LLM

reasoning using external knowledge sources to improve factual reliability. Biomedical knowledge graphs (KGs) such as PrimeKG integrate curated relationships among diseases, phenotypes, drugs, and genes, and have been leveraged to enhance factual grounding in medical question answering and decision-support systems.<sup>19,18</sup>

In this paper, we connect these previous works by integrating biomedical KG structure directly into the CoT generation process for visit-level disease prediction on structured EHR data. We propose a KG-guided CoT framework for MIMIC-III<sup>6</sup> that produces clinically grounded reasoning rather than unconstrained free-form explanations. In our formulation, each visit-level instance uses ICD-9 code features from an index visit  $t$  to predict a binary disease label at the immediately subsequent visit  $t+1$ . Concretely, we map ICD-9 codes to PrimeKG<sup>19</sup> through a three-stage entity-alignment procedure combining exact matching, embedding-based similarity retrieval, and LLM-driven clinical validation. Using these mapped entities, we identify disease-specific relevant features and extract KG paths linking each feature to its corresponding disease, pruning spurious or clinically irrelevant chains with GPT-4o. Guided by these paths, we prompt an LLM to generate temporally consistent, KG-anchored CoT explanations of whether a disease will appear at the next visit, and we retain only those reasoning traces whose conclusions match the ground-truth label. We then fine-tune lightweight open-weight models on this KG-anchored CoT supervision corpus so that, at inference time, they take ICD-9 features together with disease-specific KG evidence as input and produce both predictions and clinician-style explanations. This enables data-efficient, interpretable patient-level prognostic reasoning grounded in biomedical knowledge graphs.

We evaluate this framework on ten PrimeKG-mapped diseases in MIMIC-III across small patient cohorts (400 and 1,000 patients), where our KG-anchored CoT supervision enables LLMs to outperform strong classical baselines such as XGBoost and Random Forest, achieving AUROC of approximately 0.67 and 0.70 and macro-AUPR of approximately 0.40 and 0.47, respectively. Clinician evaluations further show that KG-guided CoT yields substantially clearer, more coherent, and more clinically sound reasoning across three dimensions: clarity & coherence, coverage & relevance, and correctness & soundness. Finally, our CoT-tuned models transfer effectively to the CRADLE cohort for forecasting cardiovascular events within one year after type 2 diabetes, improving accuracy from roughly 0.40–0.51 (untuned LLMs) to 0.72–0.77 without any retraining, and providing more temporally consistent and clinically plausible explanations. Together, these results demonstrate that KG-anchored CoT supervision enables small LLMs to learn data-efficient, clinically grounded prognostic reasoning that generalizes across cohorts.

## Method



**Figure 1:** General pipeline for KG-guided CoT data generation. **Observed clinical features** are denoted in **Blue**; **Intermediate KG nodes** are denoted in **Green**, and the **Target disease** is denoted in **Red**.

**Preliminaries.** A knowledge graph (KG) is a structured representation of entities and their relationships, typically modeled as a graph  $G = (V, E)$  where  $V$  denotes the set of entities (nodes) and  $E$  the set of typed edges encoding relations among them. In the biomedical setting, KGs organize curated knowledge about diseases, phenotypes, drugs, genes, and related concepts, and provide an explicit relational backbone that can be used to guide or constrain model reasoning. Since biomedical relations in our external knowledge source are represented as direction- and type-specific (e.g., disease  $\rightarrow$  gene and gene  $\rightarrow$  disease may be encoded as distinct edges), we treat the underlying biomedical KG as a heterogeneous *directed* multigraph. Specifically, we adopt PrimeKG<sup>19</sup> as our external knowledge source

for grounding visit-level prognostic reasoning. PrimeKG is a large-scale biomedical KG that integrates curated relationships among diseases, phenotypes, drugs, genes, and other biomedical entities into a unified graph. We use it as the substrate for (1) aligning structured EHR concepts to KG nodes, (2) identifying disease-relevant neighborhood regions, and (3) extracting reasoning paths that connect clinical features to target diseases.

**Prediction Task.** We consider a next visit disease prediction task over structured EHR data. For each patient, let  $\{x_t\}_{t=1}^T$  denote a sequence of visits ordered in time, where each visit is represented as a binary feature vector  $x_t \in \{0, 1\}^N$ , with  $N$  the dimensionality of the code space. Let  $\mathcal{D}$  denote a set of disease labels of interest, and for each  $d \in \mathcal{D}$  and visit index  $t$ , define a binary outcome  $y_{t+1}^{(d)} \in \{0, 1\}$ , indicating whether disease  $d$  is present at the next visit  $t+1$ . The model’s goal is, for each pair  $(x_t, d)$ , to estimate a probability  $\hat{p}_{t+1}^{(d)} = f(x_t, d)$  and an associated binary decision  $\hat{y}_{t+1}^{(d)}$  about whether  $d$  will appear at visit  $t+1$ . Our KG-guided models are additionally conditioned on disease-specific KG evidence (relevance sets and reasoning paths) associated with  $d$ , as described below. In our experiments, we instantiate this general setup using ICD-9 diagnosis codes as features, with  $N = 7,423$  and  $|\mathcal{D}| = 10$  diseases on the MIMIC-III cohort (details in the Results section).

**KG Entity Mapping.** To incorporate structured biomedical knowledge, we align ICD-9 concepts from MIMIC-III to PrimeKG nodes through a three-stage mapping procedure. Let  $E_{\text{ICD}}$  denote the set of unique ICD-9 descriptions, and let  $V_G$  be the node set of PrimeKG. For each entity  $e_i \in E_{\text{ICD}}$ , we retrieve a similarity-ranked candidate set  $S_i = \{s_1, \dots, s_C\} \subset V_G$  obtained by encoding  $e_i$  and all KG node labels using a text embedding model and ranking candidates by cosine similarity. The final mapped entity  $\hat{e}_i$  is selected in three stages:

**Stage 1 (Exact match).** If the ICD-9 text of  $e_i$  exactly matches a node label  $s_c$  in  $S_i$ , we record  $\tilde{e}_i = s_c$ .

**Stage 2 (Similarity match).** If an exact match is not found and the top similarity score exceeds a predefined threshold  $\tau$  (set to 0.85), we select the most similar entity:

$$\tilde{e}_i = \arg \max_{s_c \in S_i} \cos(e_i, s_c) \quad \text{if} \quad \max_{s_c \in S_i} \cos(e_i, s_c) > \tau.$$

**Stage 3 (LLM-based filtering).** All provisional mappings  $\{\tilde{e}_i\}_{i=1}^n$  from Stages 1 and 2 are then passed to GPT-4o, which validates or revises them to ensure clinical correctness, yielding the final mapped entities

$$\hat{e}_i = \text{LLM}(e_i, \tilde{e}_i \mid I_{\text{select}}),$$

where  $I_{\text{select}}$  is the prompt that asks the LLM to confirm, correct, or reject the candidate mapping based on clinical and biomedical plausibility. This three-stage procedure produces a final set of 1,513 mapped ICD-9 feature nodes and 10 mapped disease nodes, which serve as anchors for downstream relevance mining and reasoning path construction.

**KG Relevance Node and Path Mining.** Let  $\{\hat{e}_i^f\}_{i=1}^n$  denote the mapped ICD-9 feature nodes and  $\{\hat{e}_d\}_{d \in \mathcal{D}}$  denote the mapped disease nodes in PrimeKG (here  $n=1,513$  and  $|\mathcal{D}|=10$ ). Each  $\hat{e}_d$  corresponds to one of the target diseases, while the  $\hat{e}_i^f$  represent potential risk factors, comorbidities, or intermediate biomedical concepts. Our goal is to derive, for each disease  $d \in \mathcal{D}$ , a compact set of KG reasoning paths linking a small subset of feature nodes to that disease. For disease-centric node relevance, we present the disease node  $\hat{e}_d$  and the set of all mapped feature nodes  $\{\hat{e}_i^f\}_{i=1}^n$  to GPT-4o and identify the most relevant features for predicting the presence of  $\hat{e}_d$  at  $t+1$ . This yields relevance set:

$$\mathcal{R}_d = \{\hat{e}_{k,d}\}_{k=1}^{K_{\text{node}}} = \text{LLM}(\{\hat{e}_i^f\}_{i=1}^n, \hat{e}_d \mid I_{\text{node.select}}), \quad K_{\text{node}} = 8,$$

where each  $\hat{e}_{k,d}$  is a mapped feature node selected from  $\{\hat{e}_i^f\}_{i=1}^n$  for disease  $d$ . Given  $\mathcal{R}_d$ , we extract reasoning chains from PrimeKG by computing all shortest paths between each feature node  $\hat{e}_{k,d} \in \mathcal{R}_d$  and the disease node  $\hat{e}_d$ . Following prior work on KG-grounded reasoning,<sup>20,21</sup> we use shortest paths to limit “overthinking” and preserve only the most immediate biomedical relations:

$$\begin{aligned} \tilde{\mathcal{P}}_{k,d} &= \text{shortest\_path}(\hat{e}_{k,d}, \hat{e}_d, G, L), \\ \mathcal{P}_d &= \{P_{k,d}\}_{k=1}^{K_{\text{path}}} = \text{LLM}(\{\tilde{\mathcal{P}}_{k,d} : \hat{e}_{k,d} \in \mathcal{R}_d\}, \hat{e}_d \mid I_{\text{path.select}}), \quad K_{\text{path}} = 5, \end{aligned}$$

where  $\text{shortest\_path}(\cdot, \cdot; G, L)$  denotes the set of all minimum-length paths between two nodes in the biomedical knowledge graph  $G$  (PrimeKG), subject to a maximum hop bound  $L$  (we use  $L=5$  in our experiments) to avoid overly long, weakly supported chains. Since multiple paths may exist for a disease, many of which are clinically irrelevant, weakly supported, or redundant, we apply GPT-4o under a path-selection instruction  $I_{\text{path\_select}}$  to obtain a disease-level path set, yielding a compact subgraph capturing mechanistic or epidemiologic links from relevant clinical concepts to disease  $d$ . These subgraphs serve as the KG-grounded evidence used in our subsequent CoT generation.

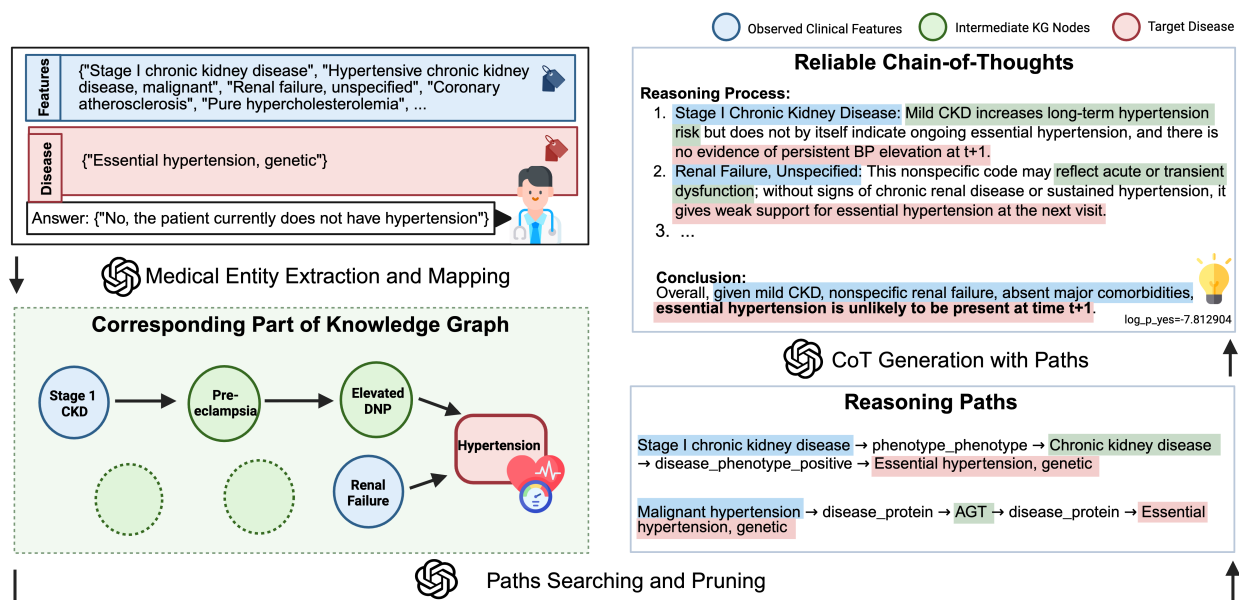
**CoT Generation with KG-Guided Reasoning.** Utilizing the disease-specific KG reasoning paths  $\mathcal{P}_d$  as guidance, we distill reliable knowledge from the off-the-shelf KG into our CoT data. To achieve this, we prompt the LLM to consider these paths together with KG-related evidence and the observed feature nodes, and to elaborate them into medically grounded CoT explanations of the prediction, represented as

$$C_{t,d} = \text{LLM}(d, x_t^+, \mathcal{R}_d^+(t), \mathcal{R}_d^-(t), \mathcal{P}_d, y_{t+1}^{(d)} | I_{\text{gen}}),$$

$$\mathcal{R}_d^+(t) = \{\hat{e}_{k,d} \in \mathcal{R}_d \mid x_t(\hat{e}_{k,d}) = 1\}, \quad \mathcal{R}_d^-(t) = \mathcal{R}_d \setminus \mathcal{R}_d^+(t).$$

Here  $I_{\text{gen}}$  is the CoT instruction prompt, and  $x_t \in \{0, 1\}^N$  is the binary ICD-9 feature vector for the index visit at time  $t$ , with  $x_t^+ = \{c \mid x_t(c) = 1\}$  denoting the set of ICD-9 codes present at that visit. The term  $y_{t+1}^{(d)} \in \{0, 1\}$  is the ground-truth label indicating whether disease  $d$  is present at the next visit  $t+1$ . The sets  $\mathcal{R}_d^+(t)$  and  $\mathcal{R}_d^-(t)$  indicate which disease-relevant KG nodes mapped from PrimeKG are expressed and not expressed at time  $t$ , respectively. The condition  $x_t(\hat{e}_{k,d}) = 1$  means that the ICD-9 code corresponding to KG node  $\hat{e}_{k,d}$  is present in the feature vector  $x_t$ .

**Filtering.** Each generated CoT  $C_{t,d}$  is required to end with an explicit binary conclusion  $\hat{y}_{t+1}^{(d)} \in \{\text{Yes}, \text{No}\}$ . We use this conclusion to decide whether to keep the example:  $C_{t,d}$  is kept if and only if the CoT-implied label matches the ground-truth outcome  $y_{t+1}^{(d)}$ , i.e.,  $\hat{y}_{t+1}^{(d)} = y_{t+1}^{(d)}$ . The surviving set of  $(x_t, d, C_{t,d}, y_{t+1}^{(d)})$  pairs forms a compact but high-quality KG-grounded supervision corpus. We fine-tune lightweight open-weight LLMs (e.g., LLaMA,<sup>22</sup> Gemma<sup>23</sup>) on this dataset so that, at test time, they produce both predictions and clinician-style reasoning traces.



**Figure 2:** Schematic and example of KG-guided CoT generation and filtering. **Observed clinical features** are denoted in **Blue**; **Intermediate KG nodes** are denoted in **Green**, and the **Target disease** is denoted in **Red**.

## Results

**Dataset.** We work with longitudinal visits from the MIMIC-III critical care database.<sup>6</sup> Following the setup of Harutyunyan et al.,<sup>24</sup> we focus on patients with more than one hospital visit and construct all pairs of adjacent visits. For each such pair, the earlier visit serves as the index visit at time  $t$  and the later visit provides labels for the next visit at time  $t+1$ . This preprocessing yields 12,353 labeled index visits (i.e., visit pairs), which we refer to as patient cases in the remainder of the paper.

Each index visit at time  $t$  is represented as a binary ICD-9 feature vector  $x_t \in \{0, 1\}^N$  with  $N = 7,423$ , where each dimension indicates the presence of a specific ICD-9 code during that visit. The original task is formulated as a multi-label prediction problem over 25 acute-care conditions derived from Clinical Classifications Software (CCS).<sup>24</sup> In this work, we adopt the same 25 disease labels but evaluate on  $|D|=10$  diseases that map cleanly to PrimeKG (Table ??). For each disease  $d \in \mathcal{D}$  and index visit  $t$ , the prediction target is a binary label  $y_{t+1}^{(d)} \in \{0, 1\}$  indicating whether  $d$  is present at the next visit  $t+1$ . Models take as input  $(x_t, d)$  (and, for certain variants, KG evidence constructed as described in the Method section) and output a probability  $\hat{p}_{t+1}^{(d)}$  and a binary verdict. We split 12,353 index visits by visit (10% test) and subsample training sets of 400 and 1,000 visits for data-efficiency experiments, using the remainder for validation.

In addition to MIMIC-III, we evaluate zero-shot transfer on Project CRADLE (Emory Clinical Research Analytics Data Lake Environment), a de-identified EHR repository from Emory Healthcare. Following prior work,<sup>25</sup> we focus on patients with type 2 diabetes and formulate a binary prediction task: whether a patient will experience a cardiovascular disease (CVD) event within one year of the initial diabetes diagnosis. CVD endpoints include coronary heart disease, congestive heart failure, myocardial infarction, and stroke, identified via ICD-9/ICD-10 codes. After applying standard inclusion/exclusion criteria from the original CRADLE studies, the resulting cohort contains 36,611 patients with 12,724 binary features.

**Evaluation Metrics.** Since both the MIMIC-III and CRADLE cohorts exhibit substantial class imbalance, we therefore report accuracy, area under the receiver operating characteristic curve (AUROC), area under the precision–recall curve (AUPR), and macro-F1 as our primary evaluation metrics. For accuracy and F1, we convert predicted probabilities to binary decisions using a fixed threshold of 0.5. Table 1 lists the 10 diseases used as prediction targets.

**Experiment:** We conduct experiments under two training regimes: one using 400 index-visit cases and another using 1,000 index-visit cases. Unless otherwise noted, all models are evaluated using a fixed decision threshold of 0.5 on the predicted probabilities.

As shown in Table 2, KG-guided CoT fine-tuning substantially boosts the performance of both LLaMA-3.1-8B and Gemma-7B in the low-data MIMIC-III setting. With only 400 labeled index visits data, the KG-CoT-tuned LLaMA model attains the best accuracy and macro-F<sub>1</sub> and the highest macro-AUPR across all methods, while matching the AUROC of the strongest logistic regression baseline ( $\approx 0.67$ ). Gemma-7B shows a similar trend: KG-CoT supervision raises its macro-AUPR from 0.25 to 0.40 and macro-F<sub>1</sub> from essentially zero to  $\approx 0.40$ , bringing it in line with the tuned LLaMA variant. When the training set is increased to 1,000 cases, both KG-guided models improve further, with LLaMA-3.1 achieving the top accuracy, AUROC ( $\approx 0.70$ ), and macro-AUPR ( $\approx 0.47$ ), and Gemma-7B achieving the highest macro-F<sub>1</sub>, consistently outperforming classical baselines such as XGBoost and Random Forest. These results indicate that KG-anchored CoT supervision enables small open-weight LLMs to learn data-efficient, clinically grounded prognostic reasoning that is competitive with or superior to traditional EHR prediction models.

**Table 1:** The 10 MIMIC-III diseases mapped into PrimeKG and used as prediction targets.

Mapped Disease Label	Category
Acute myocardial infarction	Cardiovascular
Chronic kidney disease	Renal
Chronic obstructive pulmonary disease	Respiratory
Conduction disorders	Cardiovascular
Coronary atherosclerosis	Cardiovascular
Diabetes mellitus (no complication)	Metabolic
Essential hypertension	Cardiovascular
Gastrointestinal hemorrhage	Gastrointestinal
Pneumonia	Infectious
Shock	Critical

**Table 2:** Data-efficiency comparison on MIMIC-III visit-level disease prediction with 400 and 1000 index visits (10 diseases). Best metric in **bold**, second-best underlined.

Model	MIMIC-III (400 index visits)				MIMIC-III (1000 index visits)			
	Acc	AUROC	AUPR	F1	Acc	AUROC	AUPR	F1
LLaMA3-8B (orig.)	0.4709	0.5604	0.2773	0.3327	0.4709	0.5604	0.2773	0.3327
<b>LLaMA3-8B + KG-CoT (FT)</b>	<b>0.8389</b>	<u>0.6683</u>	<b>0.4049</b>	<b>0.4050</b>	<b>0.8544</b>	<b>0.6995</b>	<b>0.4685</b>	<u>0.4058</u>
Gemma-7B (orig.)	0.7475	0.5152	0.2528	0.0169	0.7475	0.5152	0.2528	0.0169
<b>Gemma-7B + KG-CoT (FT)</b>	<u>0.8211</u>	0.6648	<u>0.4002</u>	<u>0.3976</u>	<u>0.8458</u>	0.6609	<u>0.4363</u>	<b>0.4263</b>
SGD	0.7277 ± 0.0251	0.6342 ± 0.0211	0.3595 ± 0.0235	0.1183 ± 0.0474	0.6886 ± 0.0243	0.6563 ± 0.0115	0.3840 ± 0.0114	0.1508 ± 0.0176
Logistic Regression	0.7503 ± 0.0096	<b>0.6690</b> ± 0.0054	0.3970 ± 0.0067	0.0990 ± 0.0277	0.7280 ± 0.0158	<u>0.6874</u> ± 0.0105	0.4178 ± 0.0049	0.1238 ± 0.0138
SVM	0.7385 ± 0.0000	0.6468 ± 0.0095	0.3824 ± 0.0032	0.1246 ± 0.0000	0.7391 ± 0.0011	0.6873 ± 0.0042	0.4256 ± 0.0054	0.1275 ± 0.0058
MLP	0.7334 ± 0.0105	0.6492 ± 0.0096	0.3803 ± 0.0145	0.1368 ± 0.0211	0.7029 ± 0.0299	0.6642 ± 0.0086	0.3911 ± 0.0069	0.1917 ± 0.0367
RF	0.7999 ± 0.0030	0.6155 ± 0.0084	0.3687 ± 0.0053	0.1970 ± 0.0043	0.8107 ± 0.0024	0.6574 ± 0.0076	0.4024 ± 0.0049	0.2191 ± 0.0042
Naive Bayes	0.5810 ± 0.0240	0.6010 ± 0.0162	0.3373 ± 0.0155	0.3672 ± 0.0090	0.5394 ± 0.0185	0.6224 ± 0.0039	0.3581 ± 0.0064	0.3836 ± 0.0036
XGBoost	0.8018 ± 0.0037	0.6476 ± 0.0053	0.3938 ± 0.0085	0.2090 ± 0.0058	0.8159 ± 0.0018	0.6799 ± 0.0095	0.4309 ± 0.0060	0.3065 ± 0.0058

To isolate KG guidance and filtering effects, we compare three variants at 1,000 index visits in Table 3. Removing KG guidance or using unfiltered CoT traces reduces AUROC/AUPR, suggesting that ungrounded or noisy rationales degrade ranking performance. Overall, KG-grounded and curated CoT supervision is key to strong AUROC/AUPR in this regime.

To assess cross-cohort robustness, we evaluate the MIMIC-trained models zero-shot on CRADLE (36,611 patients; 12,724 binary features) to predict cardiovascular disease within one year after type 2 diabetes diagnosis, without re-training or calibration. MIMIC-trained models transfer to the CRADLE cohort to forecast cardiovascular disease within one year after the diagnosis of type 2 diabetes. For the LLaMA-3.1 backbone, fine-tuning on 400–1,000 MIMIC-III cases increases accuracy from 0.40 to 0.75–0.77 and macro-F<sub>1</sub> from 0.34 to 0.46–0.51, while modestly raising AUROC (to ≈ 0.52–0.53) and AUPR (to ≈ 0.25). Gemma-7B exhibits a comparable pattern: accuracy rises from 0.51 to 0.72–0.77 and macro-F<sub>1</sub> from 0.31 to about 0.48, again with modest changes in AUROC and AUPR. These results suggest that the reasoning learned from KG-guided CoT supervision on MIMIC-III generalize across cohorts and spaces, enabling zero-shot transfer of backbones to out-of-distribution EHR data.

**Human Evaluation.** To evaluate whether KG-guided CoT improves explanation quality, we conducted a blinded human study comparing our KG-anchored LLaMA-3.1 model with its untuned baseline. For each of 115 randomly sampled cases, clinicians were shown an input summary, the ground-truth outcome, and two anonymized model outputs (A and B), each containing a prediction and its reasoning trace (one KG-guided, one baseline). Without knowing which model produced which trace, clinicians made a pairwise preference judgment (A vs. B) along three complementary dimensions:

1. **Clarity & Coherence (presentation):** the trace is easy to follow, logically ordered, and avoids jumps, circular reasoning, or contradictions.
2. **Coverage & Relevance (evidence selection):** the trace surfaces key patient-specific factors for the target disease, stays clinically focused, and avoids generic or distracting content.

**Table 3:** Ablation on MIMIC-III visit-level disease prediction with 1000 patient index visits. Best metric in **bold**, second-best underlined.

Setting	Acc	AUROC	AUPR	F1
LLaMA3-8B + KG (unfiltered)	0.8485	0.6589	0.4473	0.4023
LLaMA3-8B (no KG, filtered)	<b>0.8613</b>	<u>0.6640</u>	<u>0.4637</u>	<b>0.4357</b>
<b>LLaMA3-8B + KG</b>	<u>0.8544</u>	<b>0.6995</b>	<b>0.4685</b>	<u>0.4058</u>

**Table 4:** Zero-shot transfer on the CRADLE cohort. Best metric in **bold**, second-best underlined.


Model	Acc	AUROC	AUPR	F1
Llama3-8b-it (orig.)	0.4026	0.5170	0.2241	0.3446
<b>Llama3-8b-it-400.data (FT)</b>	0.7508	<u>0.5176</u>	<b>0.2512</b>	<b>0.5143</b>
<b>Llama3-8b-it-1000.data (FT)</b>	<b>0.7686</b>	<b>0.5348</b>	<u>0.2488</u>	0.4567
Gemma-7b (orig.)	0.5102	0.5052	0.2198	0.3064
<b>Gemma-7b-400.data (FT)</b>	0.7217	0.5023	0.2209	<u>0.4905</u>
<b>Gemma-7b-1000.data (FT)</b>	<u>0.7654</u>	0.4845	0.2121	0.4763

- Correctness & Soundness (clinical validity):** clinical claims and timelines are medically plausible and consistent with the provided features/outcome (no fabricated details or inappropriate causal leaps), and the conclusion follows from the stated evidence.

These dimensions are intentionally distinct: *clarity* evaluates how the reasoning is communicated, *coverage* evaluates what evidence is selected and emphasized, and *correctness* evaluates whether the medical content is accurate and logically warranted. The annotation was performed by two human experts from different medical schools: an MD professor and an MD/PhD student in biomedical informatics. In total, 115 unique cases were annotated, with the MD professor reviewing 65 cases and the MD/PhD student reviewing 50 cases.

**Inter-rater consistency.** To quantify consistency across adjudicators, we additionally double-annotated a random subset of  $n_{\text{overlap}}$  of 55 cases (both clinicians independently rated the same cases across all three dimensions) and computed Cohen’s  $\kappa$  on the binary preference decision (KG-guided vs. baseline). Agreement was  $\kappa_{\text{clar}}$  for *Clarity & Coherence*,  $\kappa_{\text{cov}}$  for *Coverage & Relevance*, and  $\kappa_{\text{corr}}$  for *Correctness & Soundness*.

On the 115 annotated cases, both clinicians preferred the KG-guided model (Model A) over the untuned baseline (Model B): 96.5% for *Clarity & Coherence*, 94.8% for *Coverage & Relevance*, and 98.3% for *Correctness & Soundness*. Qualitatively, clinicians noted that KG-guided CoT traces more often (i) organized evidence temporally, (ii) highlighted clinically meaningful risk factors over superficial code patterns, and (iii) avoided unsupported leaps or hallucinated details. By contrast, the untuned baseline was described as more “trigger-happy” (more likely to conclude *YES* from weak signals). While both models are sensitive to poor or incomplete inputs, the baseline was more prone to repetitive differentials and illogical reasoning, especially on short or noisy inputs.



**Figure 3:** Clinician preference for KG-guided explanations vs. the untuned baseline across three dimensions (donuts show % preferred; remainder is baseline).

## Discussion

We introduce a KG-guided chain-of-thought (CoT) framework that uses PrimeKG structure to scaffold LLM reasoning over structured EHRs. By mapping ICD-9 codes to PrimeKG,<sup>19</sup> mining disease-specific relevance sets and short paths, and applying label-consistency filtering, we build a compact supervision corpus for data-efficient fine-tuning of lightweight LLMs.

**Finding 1: KG-anchored CoT improves low-data learning, especially under imbalance.** Across both 400 and 1,000 labeled index visits (Table 2), KG-guided CoT substantially improves LLaMA-3.1-8B and Gemma-7B and is competitive with strong classical baselines. Gains are largest on macro-AUPR (e.g.,  $\approx 0.25 \rightarrow 0.40$  for Gemma at 400), suggesting with KG-anchored supervision encouraging disease-relevant evidence over superficial code co-occurrence.

**Finding 2: Filtering and KG evidence play complementary roles.** In the ablation (Table 3), filtering alone yields the best accuracy/ $F_1$ , while adding KG evidence yields the best AUROC (up to  $\approx 0.70$ ). Practically, filtering reduces noisy supervision and spurious positives, whereas KG structure strengthens ranking; threshold calibration can further improve accuracy/ $F_1$ . Future work can further improve threshold-dependent metrics (accuracy/ $F_1$ ) by calibrating decision thresholds per disease or cohort.

**Finding 3: Zero-shot transfer mainly shifts the operating point, motivating calibration.** On CRADLE (Table 4), fine-tuned models boost accuracy (0.72–0.77) with only modest AUROC/AUPR changes, indicating a changed decision rule under a fixed threshold rather than large ranking gains. Deployment evaluations should therefore include calibration and threshold selection, especially across cohorts with different prevalences and coding practices.

**Finding 4: Clinicians preferred KG-guided explanations.** Clinicians preferred KG-guided traces on clarity/coherence, coverage/relevance, and correctness/soundness for most cases (Section 1). They cited better temporal organization,

more clinically meaningful risk factors, and fewer unsupported leaps, suggesting KG evidence as an effective planning scaffold for clinician-facing rationales.

**Limitations and future work.** Our approach depends on the coverage and quality of PrimeKG, so gaps or biases in the KG can propagate into both supervision and explanations. The ICD to KG mapping and the relevance/path selection steps rely on GPT-4o and heuristic constraints. While we constrain candidate paths using shortest-path enumeration and use the LLM only as a selector, the pruned set is not guaranteed to be exhaustive: alternative clinically valid paths may be removed in favor of a compact scaffold. Methodologically, we also adopt a single-visit formulation ( $x_t$ ) for next-visit prediction, which simplifies evaluation but does not fully exploit longitudinal history. Future work can extend the framework to richer EHR views (notes, medications, labs), incorporate multi-visit context, and integrate clinician-in-the-loop validation for KG evidence selection and pruning.

## Conclusion

In this paper, we introduced a KG-guided chain-of-thought framework that couples biomedical structure from PrimeKG with LLM-based reasoning over structured EHR data. By using KG paths as scaffolds for CoT generation and label-consistent filtering, we constructed a compact supervision corpus that enables small open-weight LLMs to achieve competitive or superior visit-level disease prediction on MIMIC-III under strict data constraints, while also producing explicit, clinically grounded explanations. Zero-shot transfer to the CRADLE cohort suggests that the learned decision boundaries generalize across cohorts and feature spaces, and blinded clinician review indicates a consistent preference for KG-guided CoT over untuned LLM explanations. Taken together, these results highlight KG-anchored CoT supervision as a promising direction for building clinical decision-support systems in which LLMs not only predict risk but also articulate transparent, patient-level reasoning aligned with biomedical knowledge and clinician expectations.

## Acknowledgements

This research was partially supported by the internal funds and GPU servers provided by the Computer Science Department of Emory University, the US National Science Foundation under Award Numbers 2442172, 2312502, 2319449, and the US National Institutes of Health under Award Numbers K25DK135913, RF1NS139325, R01DK143456 and U18DP006922.

## References

- [1] Jennifer King, Vaishali Patel, Eric W. Jamoom, and Michael F. Furukawa. Clinical benefits of electronic health record use: national findings. *Health Services Research*, 49(1 Pt 2):392–404, 2014.
- [2] Tong Ruan, Liqi Lei, Yangming Zhou, Jie Zhai, Le Zhang, Ping He, and Ju Gao. Representation learning for clinical time series prediction tasks in electronic health records. *BMC Medical Informatics and Decision Making*, 19(Suppl 8):259, 2019.
- [3] Jiaping Zheng, Jorge Yarzebski, Balaji Polepalli Ramesh, Robert J Goldberg, and Hong Yu. Automatically detecting acute myocardial infarction events from ehr text: a preliminary study. In *AMIA Annual Symposium Proceedings*, volume 2014, page 1286, 2014.
- [4] Alexander L. Fogel and Joseph C. Kvedar. Artificial intelligence powers digital medicine. *NPJ Digital Medicine*, 1(5):1–4, 2018.
- [5] Irene Landi, Benjamin S. Glicksberg, Hyejin C. Lee, Sarah Cherng, Giovanni Landi, Mattia Danieletto, Christina Lu, Chih-lin Hsu, Ray Chen, and Joel T. Dudley. Deep representation learning of electronic health records to unlock patient stratification at scale. *NPJ Digital Medicine*, 3(1):96, 2020.
- [6] Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-Wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo A. Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3:160035, 2016.
- [7] Benjamin A. Goldstein, Ann Marie Navar, Michael J. Pencina, and John P. A. Ioannidis. Opportunities and challenges in developing risk prediction models with electronic health record data: a systematic review. *Journal of the American Medical Informatics Association*, 24(1):198–208, 2017.
- [8] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. Retain:

- An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29, 2016.
- [9] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*, pages 1721–1730, 2015.
  - [10] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
  - [11] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
  - [12] Jun Wang, Meng Fang, Ziyu Wan, Muning Wen, Jiachen Zhu, Anjie Liu, Ziqin Gong, Yan Song, Lei Chen, Lionel M. Ni, et al. Openr: An open source framework for advanced reasoning with large language models. *arXiv preprint arXiv:2410.09671*, 2024.
  - [13] Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A survey. *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.
  - [14] Yunfei Xie, Juncheng Wu, Haoqin Tu, Siwei Yang, Bingchen Zhao, Yongshuo Zong, Qiao Jin, Cihang Xie, and Yuyin Zhou. A preliminary study of o1 in medicine: Are we closer to an ai doctor?, 2024.
  - [15] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
  - [16] Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for reasoning. *Conference on Language Modeling (COLM)*, 2025.
  - [17] Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Bowen Jin, May Dongmei Wang, Joyce Ho, and Carl Yang. Ram-ehr: Retrieval augmentation meets clinical predictions on electronic health records. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 754–765, 2024.
  - [18] Liu et al. Improving large language model applications in biomedicine with retrieval-augmented generation: A systematic review, meta-analysis, and clinical development guidelines. *Journal of the American Medical Informatics Association*, 2025.
  - [19] Payal Chandak, Kexin Huang, and Marinka Zitnik. Building a knowledge graph to enable precision medicine. *Scientific Data*, 10(1):155, 2023.
  - [20] Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning. *arXiv preprint arXiv:2501.12570*, 2025.
  - [21] Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for 2+3=? on the overthinking of o1-like llms. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2024.
  - [22] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
  - [23] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
  - [24] Hrayr Harutyunyan, Hrant Khachatrian, David C Kale, Greg Ver Steeg, and Aram Galstyan. Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6(1):96, 2019.
  - [25] Ran Xu, Mohammed K. Ali, Joyce C. Ho, and Carl Yang. Hypergraph transformers for ehr-based clinical predictions. *AMIA Joint Summits on Translational Science Proceedings*, 2023:582–591, 2023.

**Implementation Details.** We fine-tuned LLaMA-3.1-8B-Instruct and Gemma-7B-Instruct on a single NVIDIA H200 (512 GB host memory; Emory HPC); full hyperparameters and training configuration are provided in our GitHub repository. We selected the checkpoint with the best macro-AUPR validation. KG-guided CoT traces were generated with GPT-4o; CoT data generation creation cost  $\approx$  \$100. Prompt templates are provided below.

### Prompt for pruning most related paths

```
messages = "you are a helpful, pattern-following medical assistant.

Given a disease to be verified and corresponding ICD features, select up to 5 most
correlated entity from relation path list based on the disease we try to verify.
### Output Format
Strictly follow the JSON structure below.
```json
{
  "Paths":
  [{"ranking": "1", "path": "sample_path_1", "reason": "reason for choosing this
path"}],
  ...
}
}
...
Disease to be verified: {disease_name}.
Paths: {
  ....
}
"
```

### Prompt for filtering most related node entities

```
messages = "you are a helpful, pattern-following medical assistant.

Given a disease to be verified and corresponding ICD features, select Top 8 most
correlated feature that are critically related to the disease.
IF there is not suitable entity in the similar entities, directly return the NONE.
### Output Format
Strictly follow the JSON structure below:
"json
{
  "selected_entity_i": {
    "name": "selected_entity_name",
    "id": "an int number, the index of the selected entity in the similar entities list, from
O to N-1
    "reason": "reason for choosing this entity"
  }
  ...
}
}
...
Disease to be verified: {disease_name}
ICD features: {
  ...
}
"
```

### Prompt for Chain-of-Thoughts Generation with paths

```
messages = [ {"role": "system", "content": "Given a disease to be verified, a list of diagnostic_features_present measured at the index
visit (time t), a subset of {potentially_relevant_exist_features} of confirmed positive features, a list of
{potentially_relevant_absent_features} of confirmed negatives, a partial set of preliminary reasoning paths, and a provided judgement
label referring to the NEXT visit (time t+1), your task is to reason step by step as if you are independently determining whether the
disease will be present at the next visit with **Yes** or **No** without prior knowledge of the given answer.
IMPORTANT TEMPORAL SETTING:
- Interpret all provided features as observations at time t (the previous/last visit).
- The target label is the presence/absence of the disease at time t+1 (the next visit).
- Prefer prognostic/persistence evidence (chronicity, structural disease, repeatedly documented diagnoses, objective abnormalities
likely to persist) over transient/acute findings.
- Do NOT assume that a disease present at time t persists to time t+1 unless there is explicit evidence of chronicity, recurrence risk, or
ongoing pathophysiology.
- *Explicitly check whether each key feature or diagnosis at time t is acute/transient or chronic/persistent. Down-weight acute findings
such as AMI, stroke, GI bleed, or sepsis unless explicit recurrence or chronic sequelae are documented.*
- **Medication caveat.** Do **not** infer disease presence from treatment context alone (e.g., vasopressors, inpatient insulin) unless
paired with a simple persistence anchor (e.g., diagnosis code, outpatient regimen, or longitudinal labs).
- **Acute default.** For acute/episodic diseases (AMI, shock, GI bleed, sepsis, stroke), predict **No** for t+1 unless there is explicit
"ongoing/recurrent/unresolved" evidence (e.g., recent transfusion plus source-localizing/therapeutic endoscopy or angiography)**.
- **Context ≠ diagnosis.** Do not treat comorbidities or typical co-medications as proof of the target disease (e.g., **CHF or nitrates
≠ CAD**) without direct evidence.
INSTRUCTIONS:
1. Explore the question and first filter out irrelevant partial reasoning paths provided for disease verification, leaving only useful paths
based on your expertise.
2. Incorporate the remaining paths and potentially_relevant_features naturally as if you discovered them yourself, then supplement
with your new reasoning path from provided features to form complete reasoning paths based on your expertise. Treat
potentially_relevant_present as confirmed present (1) and potentially_relevant_absent as confirmed absent (0) at time t.
3. Analyze and evaluate the reasoning paths with explicit attention to TEMPORAL VALIDITY, expanding on the most relevant ones
together with explicit diagnostic evidence to construct a logical, well-supported explanation. Diagnostic evidence at time t only
supports t+1 if it is chronic/persistent. Prefer objective findings (imaging, labs, pathognomonic signs) and recorded diagnoses with
known persistence over context-only signals. Distinguish between **prognostic risk** (may increase chance of disease) and
**confirmed presence** (objective diagnosis or persistent abnormality). Use risk features only as minor contextual modifiers, never as
decisive proof.
4. Base your decision ONLY on confirmed, definitive evidence (1/0). Avoid using assumptions, indirect risk factors, or therapy context
as critical evidence. Do NOT equate risk factor with a confirmed diagnosis.
5. If there are resolution cues (e.g., "unspecified hemorrhage" without source, single-episode events, acute procedures that typically
resolve), down-weight them for predicting t+1. Always check for **resolution indicators** (e.g., ""treated,"" ""status post,""
""resolved,"" ""acute episode"") and treat such findings as low-weight for t+1 prediction.
6. Do not mention the existence of predefined reasoning paths or the provided answer in your response.
7. Do not assume the given answer is correct. Instead, determine the answer solely based on your reasoning.
8. If your final conclusion contradicts the given answer, acknowledge potential discrepancies without mentioning provided answer
(e.g., ""Wait, there might be something wrong"") and refine your response again accordingly.
### Output:
Finding reasoning paths:
(you ""discover"" potential reasoning paths yourself by using the given paths if useful or generating your own if not. It should be
concise as a list of knowledge paths indexed properly.
Reasoning Process:
(Step-by-step reasoning process, do not assume the given answer is correct and do not mention the existence of answer.)
Conclusion:
(The final answer derived from your reasoning, with **Yes** or **No** printed and specified in the end.)
"""".strip()
""",
{"role": "user", "content":
"### Input:\n
f'Disease to be verified: {disease}\n\n'
f'diagnostic_features_present: {(present_features)}\n\n'
f'potentially_relevant_present_features: {(pot_rel_present)}\n\n'
f'potentially_relevant_absent_features: {(pot_rel_absent)}\n\n'
f'Reasoning Paths:\n{paths_block}\n\n'
f'Answer: {answer_yes_no}\n\n'
}
}]
```

Figure 4: Prompt templates used in our KG-guided CoT pipeline: (top-left) disease-relevant node selection, (top-right) KG path pruning and selection, (bottom) CoT generation conditioned on KG evidence and visit features.