

Zero-Shot Link Prediction in Knowledge Graphs with Large Language Models

Mingchen Li¹, Chen Ling², Rui Zhang^{1*}, Liang Zhao^{2*}

¹Division of Computational Health Sciences, Department of Surgery, University of Minnesota Twin Cities, USA

²Department of Computer Science, Emory University, USA

¹{li003378,zhan1386}@umn.edu, ²{chen.ling, liang.zhao}@emory.edu

Abstract—Zero-shot link prediction (ZSLP) on knowledge graphs aims at automatically identifying relations between given entities. Existing methods primarily employ auxiliary information to predict tail entity given head entity and its relation, yet face challenges due to the occasional unavailability of such detailed information and the inherent simplicity of predicting tail entities based on semantic similarities. Even though Large Language Models (LLMs) offer a promising solution to predict unobserved relations between the head and tail entity in a zero-shot manner, their performance is still restricted due to the inability to leverage all the (exponentially many) paths’ information between two entities, which are critical in collectively indicating their relation types. To address this, in this work, we introduce a Condensed Transition Graph Framework for Zero-Shot Link Prediction (CTLP), which encodes all the paths’ information in linear time complexity to predict unseen relations between entities, attaining both efficiency and information preservation. Specifically, we design a condensed transition graph encoder with theoretical guarantees on its coverage, expressiveness, and efficiency. It is learned by a transition graph contrastive learning strategy. Subsequently, we design a soft instruction tuning to learn and map the all-path embedding to the input of LLMs. Experimental results show that our proposed CTLP method achieves state-of-the-art performance on three standard ZSLP datasets.¹

Index Terms—Zero-Shot Link Prediction, Condensed Transition Graph, Large Language Models

I. INTRODUCTION

Knowledge graphs (KGs) [1], which are rich structured representations of real-world entities and their relations, are foundational in various applications, from semantic search to recommendation systems [2, 3, 4]. However, despite their extensive utility, most KGs suffer from an intrinsic shortcoming: incompleteness, which makes them unable to encapsulate the full breadth of evolving concepts. The task of KG link prediction therefore arises that aims to predict relation types between given entities in KGs [5]. Conventional methods [6, 7, 8] in KG link prediction learn low-dimensional representations of entities and relations, which are then used to infer links between entities. However, these methods rely on observed links to infer missing ones, limited by 1) *the data they have been trained on* and 2) *struggling to generalize to unseen entities or relations*. In light of both limitations, the motivation for Zero-Shot Link Prediction on KGs becomes apparent.

* Corresponding author

* Corresponding author

¹The code is available here: https://github.com/ToneLi/Graph_LLM_link_prediction

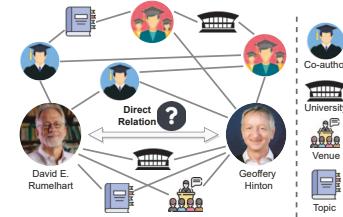


Fig. 1. Example of predicting the direct relation between two scholar entities in zero-shot by only giving their local neighboring information. Local neighboring information indicates there are many multi-hop paths with different intermediate entities, including co-author entity, venue entity, topic entity, etc.

To date, zero-shot link prediction on KGs is still under-explored, where previous works [9, 10, 11] have utilized textual features or ontologies as auxiliary information to discern relationships between head entities, seen relations, and tail entities. Specifically, they have focused on predicting tail entities given a head entity and a relation, regardless of whether that relation has been previously encountered in the training data. However, rich text features and ontologies may not always be available in KG. In addition, when given head entities and the relation, predicting tail entities is relatively easy since semantic similarity would greatly reduce the pool of candidate tail entities [12].

In this work, as illustrated in Figure 1, we focus on an exploration of a novel zero-shot link prediction task on KGs, which aims to predict unseen relation types between head and tail entities without depending on learned model parameters that tie seen relations to specific head and tail entities. As a general task solver, Large Language Models (LLMs) can intuitively be employed to predict the unseen relation [13, 14]. Existing methods [15, 16] suggest querying LLMs about the relationship between specific head and tail entities, often enriching these queries with information about neighbors. However, general neighbors are insufficient and distracting for link prediction. Instead, it is important to focus on the (higher-order) neighbors that form any paths between head and tail entities. As shown in Figure 1, to predict the direct relationship between two scholar entities: *David E. Rumelhart* and *Geoffrey E. Hinton*, there exist many multi-hop paths connecting both entities, and a high portion of these paths are related to co-authorship. This path information can provide a more concentrated source of relational cues than feeding the local neighboring information as a whole. LLMs can thus leverage path information as in-context

demonstrations to more precise predictions, i.e., co-authorship. However, extracting and summarizing these paths is fraught with computational hurdles, given the NP-hard nature of path retrieval between entities in million-scale KGs [17, 18].

To balance the trade-off between efficiency and information preservation of information encoding between head and tail entities toward link prediction, we propose a novel framework: Condensed Transition Graph Framework for Zero-Shot Link Prediction (CTLP), which distill all the paths' information into a *condensed transition graph* embedding with linear time complexity. Then, it allows LLMs to fully incorporate path information to predict relation types given a head entity and tail entity. Specifically, CTLP consists of three main components: (1) A new condensed transition graph encoder is proposed to predict the embedding aggregated from all the condensed paths between head and tail entities via different edges, with coverage, expressiveness, and efficiency guarantees; (2) A graph-centric contrastive learning is designed to learn the condensed transition graph encoder; (3) to encode the condensed graph information into the LLMs, we propose a new prefix-tuning method to encode graph information into the task instructions. We provide our key contributions as follows:

- **Problem.** We propose to solve the task of zero-shot link prediction in KGs by encoding all the paths' information in linear time complexity and letting LLMs make zero-shot predictions with soft prompt tuning.
- **Technique.** We propose to summarize all paths between given entities into a condensed transition graph while retaining the same expressiveness with a theoretical guarantee. Graph-centric contrastive learning is designed to ensure the path information is being injected into the soft prompt.
- **Experiments.** We conduct a thorough analysis of our method both quantitatively and qualitatively, which demonstrates better performance than existing state-of-the-arts. Additional experiments, including the ablation study and sensitivity analysis, illustrate the robustness of CTLP.

II. RELATED WORK

Link Prediction. Many studies [6, 11, 16, 19, 20, 21, 22] have been proposed to better predict the relationship between the head and tail entity. The translation-based model TransE [6] requires that the tail entity embedding is close to the sum of the head and relation embeddings; The non-bilinear model TuckER [20] utilizes the tucker decomposition to build the connection between different knowledge graph triples. KG-BERT [23] regards entities and relations as textual sequences, transforming the link prediction task into a classification problem. KopA [21] feeds the entity and relation embedding into LLM in the format of soft-prompting. Although performance has been achieved incrementally, these approaches in their original form are unable to learn embeddings for unseen relations. This is because they learn entities and relation embeddings using the topological structure of the KG.

Zero-shot Link Prediction. Previous works tend to predict the tail by giving the head entity and unseen relation [9]. [9] used

textual information of the relation as auxiliary information and applied a Zero-Shot Generative Adversarial Network (ZSGAN) to learn the unseen relation embedding for the task. A hierarchical n-gram framework (HNZSLP) [11] is proposed to use the structural information in relational n-grams for zero-shot link prediction. Despite the success, there exists a research gap in investigating the selection of relations from a given candidate relation set without training on the seen relation dataset. To achieve this objective, an intuitive idea is to utilize the LLMs, such as LLama [24], to predict the relations, for example, by constructing the prompt and feeding the head entity and tail entity in the format of *what is the relationship between the head entity and tail entity?* to the LLM. Unfortunately, this method is easy to cause hallucinations, because it fails to use enough knowledge to guide the LLM. So, in this work, we try to explore how to encode the path information between the head entity and the tail entity to predict the relationship in a new zero-shot learning setting.

III. PROBLEM STATEMENT

This section formulates our problem of *Zero-shot Link Prediction*.

Zero-shot link prediction on (knowledge) graphs. This task is formulated as predicting the relation type between any two entities of a knowledge graph without any training on the current knowledge graph. More specifically, we aim at predicting the relation type $r \in C_{(s,t)}$ between a head entity s and a tail entity t , where $C_{(s,t)}$ is the set of possible relation types.

The relation between s and t is well-centered among all the possible routes through which s could correlate to t , which can be collectively formulated as an (s,t) -transition graph.

a) Transition graph.: For any pair of two entities (s,t) in a knowledge graph, all the paths from head entity s to tail entity t collectively form an (s,t) -transition graph. Figure 1 exemplifies an (s,t) -transition graph where s is “David E. Rumelhart” and t is “Geoffery Hinton”. In practice, the length of paths can be upper-bounded by an integer k , which can usually be set as the diameter of the knowledge graph. To be specific, let $\mathcal{G} = (V, E)$ denote an (s,t) -transition graph consisting of a set of node V and a set of edges $E \in V \times V$. We denote by n the number of nodes in \mathcal{G} and by m its number of edges. A path from head entity s and tail entity v is denoted by $\pi = [v_0, r_0, v_1, \dots, v_{j-1}, r_i, v_j]$ such that $v_0 = s, v_j = t, (v_j, v_{j+1}) \in E$. The textual description of the path π is described as: $T(\pi) = \text{the relationship between } v_0 \text{ and } v_1 \text{ is } r_0, \dots, \text{the relationship between } v_{j-1} \text{ and } v_j \text{ is } r_i$.

IV. PROPOSED APPROACH: CTLP

A. Motivation and Approach Overview

To achieve the link prediction between s and t , the most commonly used way is to learn the embeddings of s and t with graph models (typically graph neural networks) and then calculate the similarity of their embeddings to predict their relation type r . Such a way, though efficient with $O(|E|)$, is limited in considering higher-order neighborhoods spanning

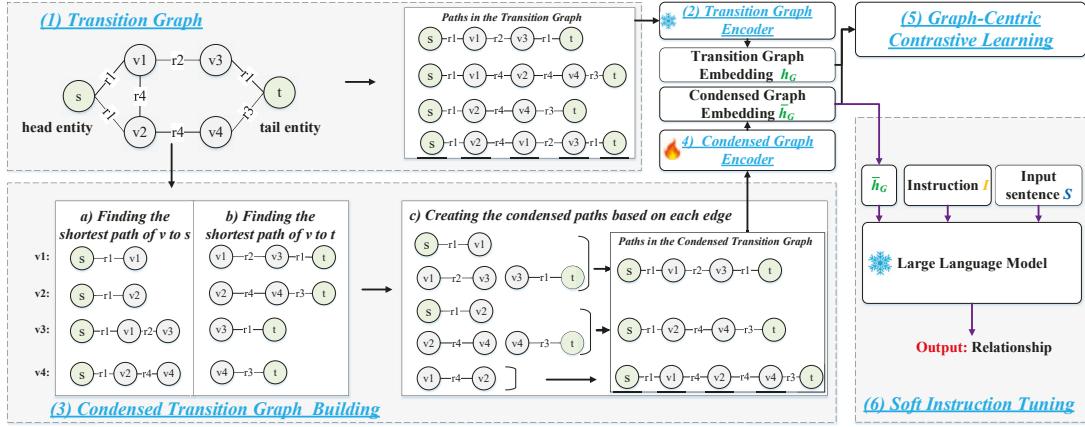


Fig. 2. Overview of CTLP. During the training process, the language model parameters are frozen, and only the condensed graph encoder is trained.

s and t due to the inherently limited number of layers of graph models. However, higher-order consideration is critical, especially in zero-shot link prediction, as it provides a broader context with more chance to generalize across different domains. On the contrary, the thorough way that preserves all the information is to consider all the possible paths traversing from s to t . However, the number of such paths is exponential to the number of edges $|E|$, which makes it prohibitive.

To address the above issue, we aim at a new framework that can not only preserve higher-hop relation information between s and t but also is efficient, with linear time complexity. More specifically, we propose a Condensed Transition Graph Enhanced LLM for the zero-shot Link Prediction (CTLP), which distills the information from all the paths between s and t to a *condensed transition graph encoder* that can learn the information in the neighborhood of all the hops with $O(|E|)$ time complexity, as detailed in Section IV-B and theoretically analyzed in Section IV-C, and illustrated in Figures 2(3) and 2(4). To let the condensed transition graph encoder better approximate the information from all the paths between s and t , we design a contrastive learning strategy that minimizes their divergence, as detailed in Section IV-D and illustrated in Figures 2(1), 2(2), and 2(5). Finally, the learned embedding will be input into LLMs with the transfrerer that is learned in Section IV-E and illustrated in Figure 2(6).

B. Condensed Transition Graph Building and Encoder

Here, the embedding of the information of each path between s and t through (u, v) can be decomposed into: (s, u) path embedding, (u, v) edge embedding, and (v, t) path embedding. To estimate the (s, u) and (v, t) path embedding, we have proposed different ways detailed in Section 4.2. We demonstrate at the end of this section that such decomposition leads to a total time complexity of $O(|E|)$.

In this section, we propose our Condensed Transition Graph encoder that can calculate the embedding in linear time complexity (as demonstrated at the end of this section). More concretely, in the condensed transition graph encoder, we leverage the principle that all the paths between s and t can

be split into each subset of paths between s and t going through each edge $(u, v) \in E$. Hence, the entirety of the embedding of all the paths between s and t is equivalent to the aggregation of the embeddings of paths pertaining to different edges $(u, v) \in E$, which is denoted as follows:

$$\bar{h}_G = \frac{1}{|E|} \sum_{(u,v) \in E} \overline{h_{(u,v)}}$$

where E is the edge set in \mathcal{G} . $\overline{h_{(u,v)}}$ denotes the embedding of all the paths between s and t through an edge (u, v) , and is calculated by the encoding of the composition of the three segments of this path, which is the key to achieving linear time complexity, as denoted in the following:

$$\overline{h_{(u,v)}} = CGE(h_{(s \rightarrow u, v \rightarrow t)})$$

where CGE is an encoder, such as the encoder part in a Transformer. $(s \rightarrow u, v \rightarrow t)$ is called a *condensed path* via an edge (u, v) and a *condensed transition graph* \mathcal{G}^* is defined as the composition of all the condensed paths via all the edges. and $h_{(s \rightarrow u, v \rightarrow t)}$ is the concatenation of the embeddings of the three segments of the path between s and t via (u, v) , namely:

$$h_{(s \rightarrow u, v \rightarrow t)} = h_{(s \rightarrow u)} | h_{(u, v)} | h_{(v \rightarrow t)}$$

where $h_{(s \rightarrow u)}$, $h_{(u, v)}$, and $h_{(v \rightarrow t)}$ are the paths' embedding from s to u , edge embedding of (u, v) , and the paths' embedding from v to t , respectively. “ $|$ ” is the concatenation function. This step is also illustrated in the section “c)” of Figure 2(3). In the following, we elaborate our efficient methods to calculate the path embedding $h_{(s \rightarrow u)}$ and $h_{(v \rightarrow t)}$, as well as edge embedding $h_{(u,v)}$, respectively.

The goal of embedding computation of the path $h_{(s \rightarrow u)}$ is to calculate the embedding that can enclose the correlation information between s and u . An intuitive option is to aggregate all the paths' embedding between s and u . However, a better trade-off between efficiency and information preservation can be a sampling or traversal through the most representative path(s) between s and t . Hence, methods like the *shortest path search* and *fattest path search* between s and t provide a reasonable trade-off. Figure 2 provides an example of this

shortest path collection, such as the shortest path of $v1$ to s is $s - r1 - v1$, its textual description $T((v1 \rightarrow s))$ is “the relationship between s and $v1$ is $r1$ ”. The embedding of $T((v1 \rightarrow s))$ is initialized by a pre-trained language model, like LLaMA. Embedding computation of the path $h_{(v \rightarrow t)}$ can be calculated similarly, while the calculation of $h_{(u,v)}$ can be done by using LLaMA to initialize its textual embedding for the statement “the relationship between u and v is r ”, r is the relationship between u and v .

C. Complexity and Expressiveness of Condensed Transition Graph Encoder

Here, we demonstrate the critical merits of our condensed transition graph encoder in terms of its coverage, expressiveness, and time complexity.

Lemma 1 (Path Coverage). *The condensed graph covers all paths from s to t when all relevant paths are within k hops.*

Proof Sketch. By finding the shortest path from s to every node v in the transition graph \mathcal{G} , we ensure that we have a way to reach every node starting from s . This guarantees that any node that is part of a path from s to t in the original subgraph is reachable in the condensed graph \mathcal{G}^* through one of these shortest paths. Similarly, by finding the shortest path from every node v to t , we ensure there is a way to reach t from any node v in \mathcal{G} . This guarantees that for any part of a path from s to t that passes through any v , there exists a continuation to t in the condensed graph \mathcal{G}^* . When all relevant paths are within k hops, for any path from s to t in \mathcal{G} , it can be decomposed into segments where each segment starts and ends at nodes where the shortest paths from s to t were calculated. Since \mathcal{G}^* contains all these shortest paths, the segments can be combined to form a path π in \mathcal{G}^* that corresponds to the original path. \square

Lemma 2 (Expressiveness). *The condensed transition graph \mathcal{G}^* retains the same expressive power of the original transition graph \mathcal{G} when all relevant paths are within k -hop in \mathcal{G} .*

Proof Sketch. First, \mathcal{G}^* ensures the connectivity from s to v for every node v in \mathcal{G}^* as proved in Lemma 1. Let $P_k(s, t)$ be the formula stating the existence of a path from s to t of length k and let $E(s, t)$ be an edge connecting the nodes s and t . $P_k(s, t)$ can be defined recursively as follows:

$$\begin{aligned} P_1(s, t) &\equiv E(s, t), \\ P_k(s, t) &\equiv \exists_v P_{k-1}(s, v) \wedge E(v, t). \end{aligned}$$

Then, \mathcal{G}^* and \mathcal{G} can achieve the WL-equivalent (i.e., \mathcal{G}^* and \mathcal{G} have the same expressive power) for $k \geq 2$ according to Theorem 3.1 and 4.1 in [25]. \square

Lemma 3 (Time complexity). *The time complexity to approximate the embedding of all the paths between s and t by our Condensed Transition Graph Encoder is linear in the number of edges in \mathcal{G} .*

Proof Sketch. The calculation of embedding consists of two steps: 1) computing the path embedding of $(s \rightarrow u)$, edge

embedding (u, v) , and path embedding $(v \rightarrow t)$, which can be done in $O(n)$, where n is the number of nodes (if using popular algorithms like breath-first search for shortest path search around s and t , respectively); and 2) aggregating the above embeddings for all the edges, which can be done in $O(m)$. Therefore, the total time complexity is $O(m)$ since \mathcal{G} is connected. \square

D. Graph-Centric Contrastive Learning

Here, we elaborate on how to train our condensed transition graph encoder proposed above. Again, our objective is to let the embedding $\bar{h}_{\mathcal{G}}$ output by the condensed transition graph encoder preserve all the information carried by all the paths in \mathcal{G} . To achieve this, we propose a graph-centric contrastive learning (GCCL) method that minimizes the information gap between them. Specifically, for each (s, t) pair, GCCL aims to reduce the distance between the embedding $\bar{h}_{\mathcal{G}}$ from the condensed transition graph encoder and the embedding $h_{\mathcal{G}}$ extracted from all the paths in \mathcal{G} (see the next paragraph about the computation of $h_{\mathcal{G}}$). To compose negative samples, we just use (s, t) for one embedding while focusing on (s', t') for the other, where $(s, t) \neq (s', t')$, by maximizing the distance between them. Note that one will tend to use small transition graph samples \mathcal{G} ’s to train the condensed transition graph encoder efficiently.

Note that the embedding of all the paths in \mathcal{G} is calculated as follows. Specifically, all the paths between the s and t in the transition graph \mathcal{G} constitute the initial paths set \mathcal{P} . For each textual path $T(\pi)$ in \mathcal{P} , its representation is learned from a pre-trained path encoder $f(\cdot)$, which can be denoted as:

$$h_{T(\pi)} = f(T(\pi)), \quad h_{\mathcal{G}} = \frac{1}{|\mathcal{P}|} \sum_{\pi}^{|\mathcal{P}|} h_{T(\pi)}$$

where the representation of \mathcal{G} (i.e., $h_{\mathcal{G}}$) is calculated by a mean function across all sample paths.

E. Soft Instruction Tuning

The embedding calculated by our condensed transition graph encoder can be used flexibly for downstream tasks. One important usage of it is a soft prompt into LLMs. To see this, here we introduce a soft instruction tuning method. In particular, we put the aligned condensed graph embedding $\bar{h}_{\mathcal{G}}$, calculated through Graph-Centric Contrastive Learning, in the front of the original instruction \mathcal{I} (*The instruction to guide the LLM for the link prediction task*) and input sentence \mathcal{S} (*what is the relationship between s and t ?*). The $\bar{h}_{\mathcal{G}}$ serves as a soft prompt in the instruction. In practice, we add several special tokens (e.g., [S]) to the beginning of the instruction, forming a sequence of length l , such as ([S] [S] [S]), where $l = 3$, and then map the representation of $\bar{h}_{\mathcal{G}}$ into these token embeddings. the concatenation of $\bar{h}_{\mathcal{G}}$, \mathcal{I} and \mathcal{S} are fed into LLM to predict the relationship between s and t . \mathcal{I} and \mathcal{S} are the embedding of \mathcal{I} and \mathcal{S} separately. In the training progress, the LLM parameter is frozen. The overall objective for our

proposed CTLP is the combination of the contrastive loss CL (elaborated in Section IV-D) and cross-entropy loss CE :

$$L_{CE} = - \sum_{i=1}^k \log(y_i | y_{1:i-1}, \bar{h}_g, \mathcal{I}, \mathcal{S})$$

$$L_{overall} = L_{CE} + L_{CL}$$

where, y_i is the generated token in the output relationship.

V. EXPERIMENTS

A. Dataset

In our experiments, we use three public KG benchmarks FB60K-NYT10 [26], UMLS [27], and NELL [28] to evaluate the capability of the proposed CTLP. FB60K-NYT10 [26] dataset includes the FB-60K knowledge graph and the NYT10 corpus. To further correct the data-building error, triples (head entity, relation, tail entity) are excluded if either the head entity or the tail entity is not found within the provided knowledge graph. UMLS [27] contains triples from the Unified Medical Language System [29], providing knowledge in the domain of healthcare and medicine. The NELL dataset, as presented by [28], consists of individual graphs corresponding to each query relation, sourced from the web. The statistic information is shown in Table I.

TABLE I
DATASETS STATISTICS, COLUMN 5 REFERS TO THE NUMBER OF TRIPLES IN THE DIFFERENT SET.

Dataset	\mathcal{V}	\mathcal{R}	# Triples	# Train/ # Dev/ # Test
FB60K-NYT10	9,840	56	13,837	12,104/-1,733
UMLS	135	46	6,529	5,216/652/661
NELL	9,093	12	12,108	8,747/543/2,818

B. Baselines and Evaluation Metrics

In our experiments, the baselines include three commonly used KG embedding methods: TransE [6], DisMult [19] and ComplEx [30]. Obviously, these original models cannot predict the unseen relationship in the zero-shot setting. Therefore, based on these three methods, we propose three zero-shot baselines, **1) ZSL-TransE**, **2) ZSL-DistMult**, and **3) ZSL-ComplEx**. Specifically, the pre-trained BERT model [31] is employed to compute embeddings for each entity and relation. Following this, the score for each candidate relation is determined using the scoring functions within TransE, DistMult, and ComplEx. **4) ZS-BERT** [32], it is designed specifically for zero-shot text classification, we adapt it to our zero-shot setting by calculating the similarity between the sentence *what is the relationship between head entity and tail entity?* and the candidate relations. Otherwise, We also compare the performance of CTLP with several strong baselines based on the LLM, including **5) GPT-3.5/4**: In GPT-3.5/4, we design hard prompts to guide the GPT models in predicting relations between the head and tail entities. **6) ZSL-InstructGLM** [16]: we employ the same prompt construction method as InstructGLM [16] to predict the relationship between two entities in zero-shot settings. We consider **7) LLAMA family** as the baselines: LLaMA2-(13, 70b) [24]. We report the standard Micro Precision, Recall, and F1-score on the test set.

C. Implementation Detail

To ensure the zero-shot setting, in the testing progress of dataset FB60K-NYT10, and UMLS, the condensed graph encoder is trained in the knowledge graph NELL, and the LLM is frozen in the training progress of the condensed graph encoder. Following this, the trained condensed graph encoder is employed to encode the condensed transition graph in FB60K-NYT10 and UMLS. The condensed graph encoder in the NELL dataset is trained using the FB60K-NYT10. We set the hops $k = 4$ in the transition graph of FB60K-NYT10, UMLS, and NELL. During the training process of the condensed graph encoder, the vector dimensions of the transition graph and the LLM are the same. In the dataset FB60K-NYT10, the length of the soft prompt token is set to $l = 5$, while in the dataset UMLS and NELL, the length of the soft prompt token is set to $l = 10$. The relation types of these three data sets do not overlap with each other.

D. Main Results

Table II presents the experiment results of various approaches based on Precision, Recall and F1. We have the following observations: (1) our CTLP significantly outperforms all the strong baselines across all evaluation metrics. (2) In the dataset FB60K-NYT10, we observe that CTLP improve the original LLaMA2 13B, LLaMA2 70B by 6.87%, 1.2% respectively, in terms of F1. (3) CTLP does not significantly enhance the performance of LLM in UMLS. This can be attributed to the ongoing challenges faced by large language models (LLMs) in understanding path information within the biomedical domain, particularly due to the complexity of relations and entity names, such as abbreviations. (4) We observe that ZSL-InstructGLM gets zero performance in the UMLS, the reason is that ZSL-InstructGLM injects the hop3 information around the head and tail entity to the LLM, and this hard prompt design exceeds the maximum input length of LLM. (5) GPT-3.5/4 exhibits the lowest performance. The main reason is that the reported results are in the zero-shot setting due to the unavailability of open resources.

VI. ANALYSIS

In order to further explore the effectiveness of our framework, we perform a series of analyses based on different characteristics of our model. First, we evaluate the effectiveness of our proposed CTLP with the models that feed all paths from the transition graph to the LLM. The contribution of our model components can also be learned from ablated models. So we propose two model variants to help us validate the advantages of the contrastive learning operation and condensed graph encoder. Next, we explore the performance of our model with a different number of hops k in the transition graph, and we also explore the model performance with different soft prompt token lengths l . Lastly, to further investigate the effectiveness of our model, we explore the model performance by using the limited dataset to train our graph encoder. Due to space constraints, please refer to Appendix IX-B and Appendix IX-C for the parameter

TABLE II
RESULTS OF VARIOUS APPROACHES FOR ZERO-SHOT LINK PREDICTION ON THREE OPEN DATASETS.

Approach	FB60K-NYT10			UMLS			NELL		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
ZSL-TransE	7.27	7.27	7.27	3.78	3.78	3.78	8.94	8.94	8.94
ZSL-ComplEx	1.21	1.21	1.21	1.81	1.81	1.81	9.44	9.44	9.44
ZSL-DistMult	7.09	7.09	7.09	4.08	4.08	4.08	7.94	7.94	7.94
ZS-BERT [32]	3.52	3.52	3.52	2.11	2.11	2.11	0.18	0.18	0.18
ZSL-InstructGLM [16]	0.52	0.52	0.52	0.00	0.00	0.00	0.19	0.19	0.19
GPT-3.5	20.00	20.00	20.00	1.98	1.98	1.98	35.29	36.00	35.64
GPT-4	12.00	12.00	12.00	6.00	6.00	6.00	39.00	39.00	39.00
LLaMA2 13b	36.18	36.18	36.18	9.83	9.83	9.83	38.18	38.18	38.18
LLaMA2 13b + CTLP	43.05	43.05	43.05	10.30	10.30	10.30	39.32	39.32	39.32
LLaMA2 70b	36.75	36.75	36.75	13.76	13.76	13.76	55.28	55.28	55.28
LLaMA2 70b + CTLP	37.97	37.97	37.97	13.92	13.92	13.92	56.78	56.78	56.78

exploration and evaluation of model performance using the limited dataset.

A. Comparison with All Paths Input

TABLE III
F1 RESULTS OF CTLP, PATHLLM-HOP-K AND CPATHLLM-HOP-K

	Approach	FB60K-NYT10	UMLS	NELL
LLaMA2 13b	PATHLLM-HOP-3	21.29	0.00	12.67
	PATHLLM-HOP-4	4.22	0.00	3.10
	CPATHLLM-HOP-3	21.29	0.00	12.67
	CPATHLLM-HOP-4	4.33	0.00	3.21
	CTLP	43.05	10.30	39.32
LLaMA2 70b	PATHLLM-HOP-3	30.92	0.00	18.29
	PATHLLM-HOP-4	6.38	0.00	6.39
	CPATHLLM-HOP-3	30.92	0.00	18.29
	CPATHLLM-HOP-4	6.26	0.00	6.42
	CTLP	37.97	13.92	56.78

To further evaluate the effectiveness of our proposed CTLP, we introduce PATHLLM-HOP-K and CPATHLLM-HOP-K, two LLMs that utilize path information to predict the relationship between the head and tail entities. In PATHLLM-HOP-K, the path information is derived from the transition graph, whereas in CPATHLLM-HOP-K, the path information is from the condensed transition graph. More specifically, we first transfer all paths in the transition graph or condensed transition graph into their textual descriptions. For example, the path $s - r1 - v1 - r2 - t$ is represented by *the relationship between v1 and s is r1, the relationship between v1 and t is r2*. Secondly, these path descriptions are fed into the LLM to predict the relationship between the head entity and the tail entity in the format of the hard prompt.

Table III shows the performance comparison. Our model demonstrates superior performance on these datasets, outperforming the PATHLLM-HOP-K and CPATHLLM-HOP-K. This is because when we provide all path information to the LLM, the prompt length exceeds the maximum input length of LLM, particularly evident in the UMLS dataset. As a result, PATHLLM-HOP-K achieves zero performance in UMLS. Otherwise, despite our algorithm could reduce the number of paths, the condensed paths still surpass the maximum input

length of the LLM. Consequently, CPATHLLM-HOP-K also exhibits a worse performance. Moreover, inputting excessively long textual descriptions to the LLM can also lead to a reduction in generation time. In contrast, our method employs a soft prompt strategy to alleviate this issue. Additionally, it utilizes contrastive learning to ensure comprehensive path information. For the comparison of path numbers on the transition graph and condensed transition graph, please check Appendix IX-A.

B. Ablation Experiments

We introduce two ablated models of CTLP: (1) **CTLP-WCL** does not use the pre-computed transition graph information to guide the learning progress of condensed graph encoder; (2) **CTLP-GCN** uses the traditional Graph Convolutional Network (GCN) [33] to encoder the condensed graph information. To ensure a fair comparison, we maintain consistency in the LLM, hop-k, and the length of prefix soft token across all datasets. We find that the performance of CTLP degrades as we remove important model components. Specifically, both CTLP-WCL and CTLP-GCN perform poorly when compared to CTLP, indicating the importance of using the pre-computed transition graph information to guide the condensed graph encoder, thereby preventing the loss of valuable path information. Otherwise, the lower performance of CTLP-GCN also indicates that GCN may face challenges in effectively learning comprehensive path information.

TABLE IV
CTLP PERFORMANCE AND ITS ABLATED MODEL ON THE F1 SCORE.

Approach	FB60K-NYT10	UMLS	NELL
CTLP-WCL	39.12	13.46	55.93
CTLP-GCN	35.49	13.61	55.78
CTLP	43.05	13.92	56.78

VII. CONCLUSION

In this paper, we introduce CTLP, a novel ZSL framework for link prediction. Our approach focuses on leveraging all sample paths between the head and tail entities to predict their relationship. To achieve this, we develop a condensed transition graph construction method and employ contrastive learning to balance time efficiency and comprehensiveness when encoding

these paths. Subsequently, the learned condensed transition graph is used as the soft prompt to feed into the LLM. Experimental results show that our framework achieves consistent improvements over various baselines in three open datasets.

VIII. ACKNOWLEDGMENT

This work was supported by the National Institutes of Health's National Center for Complementary and Integrative Health under grant number R01AT009457, National Institute on Aging under grant number R01AG078154 and National Cancer Institute under grant number R01CA287413. The content is solely the responsibility of the authors and does not represent the official views of the National Institutes of Health.

This work was supported by the National Science Foundation (NSF) Grant No. 2403312, No. 1841520, No. 2007716, No. 2007976, No. 1942594, No. 1907805, and Amazon Research Award.

REFERENCES

- [1] M. Li, Z. Zhou, and Y. Wang, “Multi-fusion chinese wordnet (mcw): Compound of machine learning and manual correction,” *arXiv preprint arXiv:2002.01761*, 2020.
- [2] M. Li and S. Ji, “Semantic structure based query graph prediction for question answering over knowledge graph,” *arXiv preprint arXiv:2204.10194*, 2022.
- [3] M. Li and L. Huang, “Understand the dynamic world: An end-to-end knowledge informed framework for open domain entity state tracking,” *arXiv preprint arXiv:2304.13854*, 2023.
- [4] M. Li, M. Chen, H. Zhou, and R. Zhang, “PetaIor: Improving large language model by tailored chunk scorer in biomedical triple extraction,” *arXiv preprint arXiv:2310.18463*, 2023.
- [5] M. Wang, L. Qiu, and X. Wang, “A survey on knowledge graph embeddings for link prediction,” *Symmetry*, vol. 13, no. 3, p. 485, 2021.
- [6] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, “Translating embeddings for modeling multi-relational data,” *Advances in neural information processing systems*, vol. 26, 2013.
- [7] S. M. Kazemi and D. Poole, “Simple embedding for link prediction in knowledge graphs,” *Advances in neural information processing systems*, vol. 31, 2018.
- [8] A. Rossi, D. Barbosa, D. Firmani, A. Matinata, and P. Merialdo, “Knowledge graph embedding for link prediction: A comparative analysis,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 15, no. 2, pp. 1–49, 2021.
- [9] P. Qin, X. Wang, W. Chen, C. Zhang, W. Xu, and W. Y. Wang, “Generative adversarial zero-shot relational learning for knowledge graphs,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8673–8680.
- [10] Y. Geng, J. Chen, Z. Chen, J. Z. Pan, Z. Ye, Z. Yuan, Y. Jia, and H. Chen, “Ontozsl: Ontology-enhanced zero-shot learning,” in *Proceedings of the Web Conference 2021*, 2021, pp. 3325–3336.
- [11] M. Li, J. Chen, S. Mensah, N. Aletras, X. Yang, and Y. Ye, “A hierarchical n-gram framework for zero-shot link prediction,” *arXiv preprint arXiv:2204.10293*, 2022.
- [12] W. Zeng, X. Zhao, W. Wang, J. Tang, and Z. Tan, “Degree-aware alignment for entities in tail,” in *Proceedings of the ACM SIGIR*, 2020, pp. 811–820.
- [13] C. Ling, X. Zhao, J. Lu, C. Deng, C. Zheng, J. Wang, T. Chowdhury, Y. Li, H. Cui, T. Zhao *et al.*, “Domain specialization as the key to make large language models disruptive: A comprehensive survey,” *arXiv preprint arXiv:2305.18703*, vol. 2305, 2023.
- [14] G. Bai, Z. Chai, C. Ling, S. Wang, J. Lu, N. Zhang, T. Shi, Z. Yu, M. Zhu, Y. Zhang *et al.*, “Beyond efficiency: A systematic survey of resource-efficient large language models,” *arXiv preprint arXiv:2401.00625*, 2024.
- [15] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, “Unifying large language models and knowledge graphs: A roadmap,” *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [16] R. Ye, C. Zhang, R. Wang, S. Xu, and Y. Zhang, “Natural language is all a graph needs,” *arXiv preprint arXiv:2308.07134*, 2023.
- [17] G. Michel, G. Nikolenzos, J. F. Lutzeyer, and M. Vazirgiannis, “Path neural networks: Expressive and accurate graph neural networks,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 24737–24755.
- [18] C. Ling, X. Zhang, X. Zhao, Y. Liu, W. Cheng, M. Oishi, T. Osaki, K. Matsuda, H. Chen, and L. Zhao, “Open-ended commonsense reasoning with unrestricted answer candidates,” in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 8035–8047.
- [19] B. Yang, W.-t. Yih, X. He, J. Gao, and L. Deng, “Embedding entities and relations for learning and inference in knowledge bases,” *arXiv preprint arXiv:1412.6575*, 2014.
- [20] I. Balažević, C. Allen, and T. M. Hospedales, “Tucker: Tensor factorization for knowledge graph completion,” *arXiv preprint arXiv:1901.09590*, 2019.
- [21] Y. Zhang, Z. Chen, W. Zhang, and H. Chen, “Making large language models perform better in knowledge graph completion,” *arXiv preprint arXiv:2310.06671*, 2023.
- [22] L. Yao, J. Peng, C. Mao, and Y. Luo, “Exploring large language models for knowledge graph completion,” *arXiv preprint arXiv:2308.13916*, 2023.
- [23] L. Yao, C. Mao, and Y. Luo, “Kg-bert: Bert for knowledge graph completion,” *arXiv preprint arXiv:1909.03193*, 2019.
- [24] H. Touvron, T. Lavigne, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [25] C. Graziani, T. Drucks, M. Bianchini, T. Gärtner *et al.*, “No pain no gain: More expressive gnns with paths,” in *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*, 2023.

[26] C. Fu, T. Chen, M. Qu, W. Jin, and X. Ren, “Collaborative policy learning for open knowledge graph reasoning,” *arXiv preprint arXiv:1909.00230*, 2019.

[27] S. Kok and P. Domingos, “Statistical predicate invention,” in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 433–440.

[28] R. Das, S. Dhuliawala, M. Zaheer, L. Vilnis, I. Durugkar, A. Krishnamurthy, A. Smola, and A. McCallum, “Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning,” *arXiv preprint arXiv:1711.05851*, 2017.

[29] D. A. Lindberg, B. L. Humphreys, and A. T. McCray, “The unified medical language system,” *Yearbook of medical informatics*, vol. 2, no. 01, pp. 41–51, 1993.

[30] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard, “Complex embeddings for simple link prediction,” in *International conference on machine learning*. PMLR, 2016, pp. 2071–2080.

[31] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” *arXiv preprint arXiv:1908.10084*, 2019.

[32] C.-Y. Chen and C.-T. Li, “Zs-bert: Towards zero-shot relation extraction with attribute representation learning,” *arXiv preprint arXiv:2104.04697*, 2021.

[33] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *arXiv preprint arXiv:1609.02907*, 2016.

IX. APPENDIX

A. Path Numbers Comparison

Table V presents a comparison of path numbers between the transition graph and the condensed transition graph. The results demonstrate that our method effectively reduces the path numbers on the transition graph.

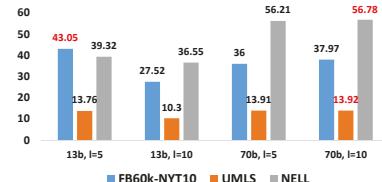
TABLE V

A COMPARISON OF PATH NUMBERS WITHIN HOP 4 IN TRANSITION GRAPH AND CONDENSED TRANSITION GRAPH. EACH VALUE REFERS TO THE AVERAGE PATH NUMBERS ON THE TEST SET ACROSS THE RELEVANT DATASET.

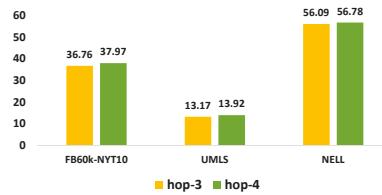
	FB60K-NYT10	UMLS	NELL
path numbers in transition graph	4.689	12.617	1.220
path numbers in condensed transition graph	1.748	558	524

B. Impact of Hop-k and the Length l of Soft Prompt Tokens

In our proposed model, there is one parameter controlling the size of the transition graph, it is hop- k . And the other important parameter is the length l of the soft prompt tokens in CTLP, we treat these two parameters with the same importance. From Figure 3(a), we observed that the model performs better under hop-4 than under hop-3. This indicates that richer path information is beneficial for relationship prediction. From Figure 3(b), For the UMLS and NELL datasets, our model demonstrates superior performance when $l = 10$, whereas its effectiveness is more pronounced on the FB60k-NYT10 dataset when $l = 5$.



(a)



(b)

Fig. 3. (a): CTLP performance with different prompt token length (l) on the F1 score. We fix the hop- k = 4. 13b denotes the LLaMA2 13b, 70b denotes the LLaMA2 70b. (b): CTLP performance with different hop- k on the F1 score. We fix the LLM as LLaMA 70b, and the $l = 10$.

C. Train the Condensed Graph Encoder using a Limited Dataset

In our study, we initially trained the condensed graph encoder on the knowledge graph NELL. Subsequently, we employ the trained condensed graph encoder to encode the condensed transition graphs in the UMLS and FB60K-NYT10 datasets. This approach ensures compatibility with the zero-shot setting. In this section, we mainly focus on investigating the effectiveness of our model when utilizing a constrained dataset for training the condensed graph encoder. For example, when training the condensed graph encoder on the knowledge graph NELL, we mask 30/50 of the relation types along with their associated entity pairs in the training dataset. The selection is based on the relation set identified in the test dataset. They are denoted as CTLP MASK 30% and CTLP MASK 50%.

The results are presented in Table VI. It is noteworthy that our model consistently outperforms the original LLM even when the mask operation is applied. The model performance even beyond the CTLP, indicating that our model achieves superior performance with a limited dataset.

TABLE VI
CTLP PERFORMANCE WITH DIFFERENT SETTINGS. FOR THE FB60K-NYT10 DATASET, THESE MODELS ARE BUILT UPON LLaMA2 13B, WHEREAS FOR THE UMLS AND NELL DATASETS, THESE MODELS ARE BUILT UPON LLaMA2 70B.

Approach	FB60K-NYT10	UMLS	NELL
LLAMA2 13B	36.18	–	–
LLAMA2 70B	–	13.76	55.28
CTLP MASK 50%	36.98	13.92	56.85
CTLP MASK 30%	42.35	14.22	56.42
CTLP	43.05	13.92	56.78