
Natural Language is All a Graph Needs

Ruosong Ye
Rutgers University
ruosong.ye@rutgers.edu

Caiqi Zhang
University of Cambridge
cz391@cam.ac.uk

Runhui Wang
Rutgers University
runhui.wang@rutgers.edu

Shuyuan Xu
Rutgers University
shuyuan.xu@rutgers.edu

Yongfeng Zhang
Rutgers University
yongfeng.zhang@rutgers.edu

Abstract

The emergence of large-scale pre-trained language models, such as ChatGPT, has revolutionized various research fields in artificial intelligence. Transformer-based large language models (LLMs) have gradually replaced CNNs and RNNs to unify fields of computer vision and natural language processing. Compared with the data that exists relatively independently such as images, videos or texts, graph is a type of data that contains rich structural and relational information. Meanwhile, natural language, as one of the most expressive mediums, excels in describing complex structures. However, existing work on incorporating graph learning problems into the generative language modeling framework remains very limited. As the importance of LLMs continues to grow, it becomes essential to explore whether LLMs can also replace GNNs as the foundation model for graphs. In this paper, we propose **InstructGLM (Instruction-finetuned Graph Language Model)**, systematically design highly scalable prompts based on natural language instructions, and use natural language to describe the geometric structure and node features of the graph for instruction tuning an LLM to perform learning and inference on graphs in a generative manner. Our method exceeds all competitive GNN baselines on ogbn-arxiv, Cora and PubMed datasets, which demonstrates the effectiveness of our method and sheds light on generative large language models as the foundation model for graph machine learning. The code and data are available at <https://github.com/agiresearch/InstructGLM>.

1 Introduction

Before the advent of Transformers [1], various artificial intelligence domains with different inductive biases had diverse foundational model architectures. For instance, CNNs [2, 3] were designed with considerations for spatial invariance in images, leading to superior performance in computer vision tasks [4, 5]. Memory-enhanced models like RNNs [6] and LSTM [7, 8] were widely used for handling sequential data such as natural language [9] and audio [10]. Graph Neural Networks (GNNs) excel in capturing topological information by employing message passing and aggregation mechanisms, making them a preferred choice in the field of graph learning for a long time [11–13].

In recent years, the AI community has witnessed the emergence of numerous powerful pre-trained Large Language Models (LLMs) [14–18], which are driving huge advancements and lead to the pursuit of possible Artificial General Intelligence (AGI) [19]. Under this background, there is a trend towards unification in model architectures across different domains. Specifically, pre-trained Transformers have demonstrated remarkable performance on various modalities, such as images [20] and videos [21] in computer vision, text in natural language processing [22], structured data in graph machine learning [23], personalized data in recommender systems [24], decision sequences in reinforcement learning [25], and visual-text pairs in multimodal tasks [26]. There has even been Transformers capable of handling twelve modalities [27].

Besides model architecture, the unification of processing method in handling multimodal data is also a significant trend worth attention. T5 [15] established a text-to-text framework, unifying all NLP

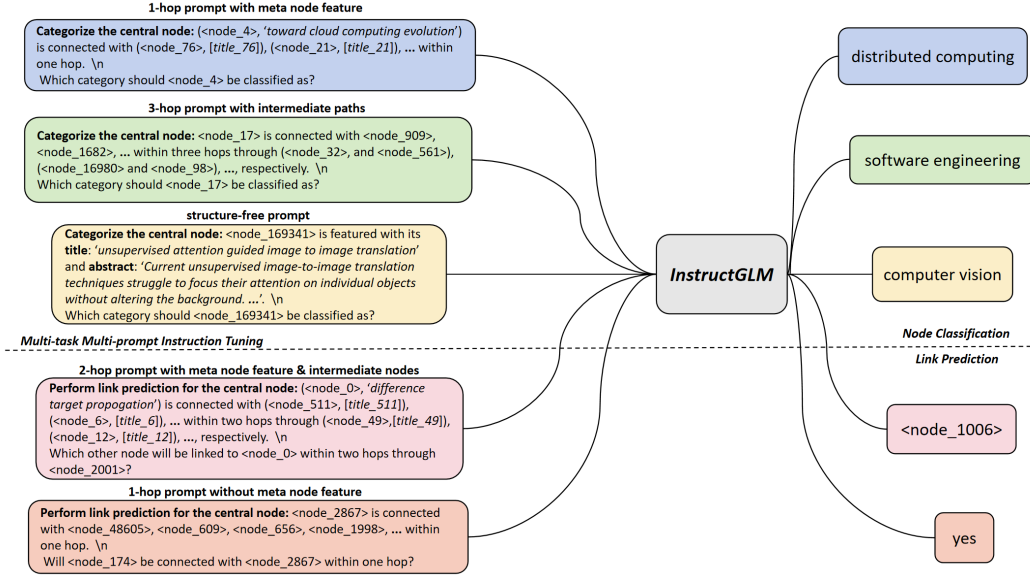


Figure 1: Illustration of the InstructGLM Framework. We fine-tune InstructGLM under a Multi-task Multi-prompt instruction tuning framework, enabling it to solve various graph machine learning tasks with the structure information purely described by natural language.

tasks as a sequence generation problem. Moreover, models like CLIP [26] utilize image-text pairs to accomplish multimodal tasks with the images captioned by natural language. As for reinforcement learning, Di Palo et al. [25] employs natural language to describe environmental states for the agent which successfully solves many reinforcement learning (RL) problems. P5 [24] further contributes to this trend by reformulating all personalized recommendation tasks as language modeling tasks via prompts. The aforementioned works collectively demonstrate that employing natural language for multimodal information representation has emerged as a prominent and promising trend.

However, in graph machine learning, the exploration for using natural language to handle graph-related tasks still remains limited. Existing methods that utilize large language models for graph tasks can be roughly categorized into two types: 1) Combining LLMs and GNNs, where the LLM acts as a feature extractor or data augmentation module to enhance the downstream GNNs’ performance [28–30]. Such kind of methods often require training multiple models and thus incurring significant computational overhead. Furthermore, since the GNNs are still responsible for learning the graph’s structural information, they tend to inherit the drawbacks of GNNs such as over-smoothing easily [31]. 2) Only relying on Transformers but necessitating novel designs of token embedding for nodes and edges [32] or creating complex graph attention module to learn structural information [33, 34]. This type of method demands local attention calculation on every node during each optimization step, leading to considerable computation costs and thus limiting each node’s scope to only 1-hop neighbors. Meanwhile, the complex pipeline conveys structural information through special attention mechanisms or token representations, which prevents the model from directly observing and learning structural information like GNNs, thus restricting further improvement on performance.

To address the issues present in LLM-based graph learners and bridge the gap of natural language based graph learning, we propose **InstructGLM (Instruction-finetuned Graph Language Model)**. Given that LLMs have been dominant in many AI domains, we aim to answer the question: **Can LLMs also replace GNNs as the foundation model in graph machine learning?** Intuitively, as one of the most expressive medium, natural language is adept at describing complex structures such that InstructGLM owns following advantages over GNNs:

- 1) *Flexibility*. A natural language sentence is capable of effectively describing the connectivity at any desired hop levels and intermediate paths without message passing and aggregation. Even the multimodal features of the nodes and edges can be directly integrated into natural language prompts, making natural language a very flexible medium to convey both structural and content information on the graph.

- 2) *Scalability*. Injecting graph structure into multiple natural language sentences enables mini-batch training and independent gradient propagation, which further allows easy scalability to distributed training and inference on massive graphs with low machine communication overhead.
- 3) *Compatibility*. Aided by structure descriptions, InstructGLM can consistently reformulate various graph learning pipelines as language modeling tasks, thus fits well into the LLM-based multimodal processing framework, paving the way to integrate graph learning with other AI tasks such as vision, language and recommendation to construct unified AI systems.

In this paper, we focus on tackling the graph node classification task, while augmenting it with self-supervised link prediction to enhance the performance. Inspired by various message passing pipelines in GNNs [35, 36], we design a series of scalable graph prompts for instruction tuning on generative LLMs [37, 38]. Specifically, after determining the central node and conducting neighbor sampling, we systematically employ natural language to describe the graph’s topology according to the prompts. By doing so, the graph structure is clearly and intuitively provided to LLMs without complex pipelines tailored to graphs. Therefore, we can handle graph tasks efficiently and succinctly by the vanilla Transformer architecture [1] and language modeling objective [39] in a generative manner. Moreover, our approach ensures high compatibility between graph learning, NLP, as well as multimodal processing, thus exhibits high scalability for multitask learning across various domains. Overall, our contributions can be summarized by the following four points:

- To the best of our knowledge, we are the first to propose purely using natural language for graph structure representation and perform instruction tuning on a generative LLM to solve graph-related problems. We eliminate the requirement of designing specific complex attention mechanisms or tokenizers tailored for graphs. Instead, we offer a concise and efficient natural language processing interface for graph machine learning, which exhibits high scalability to a unified multimodal and multitask framework, aligning with the current trend in other AI domains.
- Inspired by various message passing mechanisms in GNNs, we have designed a series of rule-based, highly scalable instruction prompts for general graph structure representation and graph machine learning. Although in this paper, our focus lies in exploring instruction tuning on large language models, these prompts can also be used for zero-shot experiments on LLMs.
- In addition to node classification, we also conduct self-supervised link prediction as an auxiliary task and further investigate its influence on the primary task under a multitask instruction tuning framework. This exploration holds valuable insights for future LLM-based multitask graph learning, demonstrating the significance of self-supervised link prediction for large language models’ better structure understanding on graphs.
- We implement extensive experiments on three widely used datasets: ogbn-arxiv, Cora, and PubMed. The results demonstrate our InstructGLM outperforms previous competitive GNN baselines and Transformer-based methods across all three datasets, achieving the top-ranked performance. These findings validate the effectiveness of our proposed method and underscore the trend of leveraging generative large language models as the foundation model for graph machine learning.

2 Related Work

2.1 GNN-based Methods

Graph Neural Networks (GNNs) [40, 41] have been dominant in the field of graph machine learning for a long period. Leveraging message passing and aggregation mechanisms, GNNs excel in simultaneously learning node features, edge features, and topological structure information. Overall, GNNs can be categorized as spatial-based ones [12, 13, 42, 43] and spectral-based ones [11, 44, 45], with various message passing mechanisms. To address some inherent problems such as over-smoothing [31], researchers have proposed methods such as merging intermediate layer features into the final representation [36], conducting convolutions on multiple subgraphs extracted from varying hop levels [35], and dropping edges to prevent overfitting [46]. One major drawback of GNNs is their inability to directly process raw data from various modalities, requiring extensive feature engineering as a preprocessing step. GNNs cannot directly handle non-numeric information such as text or images. To handle this issue, existing works use techniques such as BoW, TF-IDF, or Skip-gram to construct shallow embeddings as input to GNNs [47]. Its lack of compatibility with existing large-scale generative models presents a significant challenge for integration with other AI domains such as vision and language into a unified intelligent system.

2.2 Transformer-based Methods

Attention-based Transformer models can also be utilized for graph processing by representing each node and edge in the graph as distinct tokens [48]. However, this simple approach presents two challenges: Firstly, it becomes computationally intensive for handling large-scale graphs. Secondly, the global weighted average calculation with basic attention mechanisms can not effectively capture and learn the graph’s topological structure [32]. To overcome these issues, various approaches have been proposed to improve Transformer structures or graph representation methods. Some methods incorporate graph structure information into attention matrices [23] or coefficients [49], while others restrict attention to local subgraphs [34] or ingeniously design orthogonal vectors for node and edge tokens to encode structural details [32]. These enhancements often involve complex attention mechanisms or data transformations, making direct representation of graph structure challenging and significantly increasing the difficulty of model training. The only work similar to ours is Zhang et al. [50], which utilizes an encoder-only model and natural language templates specially formulated to solve biological concept linking problems [51, 52]. However, unlike our approach, it is not designed for general graph learning and difficult to be extended beyond the classification task due to the use of encoder-only model [53]. Additionally, its natural language templates are tailored for the domain of biological concept linking and thus not as expressive and flexible as our approach.

2.3 Fuse GNN and Transformers

GNNs excel at learning structural information, while Transformers are proficient in capturing multi-modality features. Many works have combined GNNs and Transformers to effectively address graph-related tasks. For example, Chien et al. [54] utilizes the multi-label neighbor prediction task to incorporate structural information into language models, generating enhanced features named GIANT to improve downstream GNN’s performance. Mavromatis et al. [29] employs GNNs to perform knowledge distillation on language models, Zhao et al. [30] trains GNNs and language models iteratively in a variational inference framework, and Rong et al. [55] attempts to replace the attention heads in Transformers with GNNs to better capture global information. The main drawback of the aforementioned approaches is the lack of decoupling between Transformer models and GNNs, which requires training multiple models and easily incurs significant computational overhead [34]. Moreover, the model performance is still susceptible to inherent issues of GNNs, such as over-smoothing [56]. Additionally, the pipeline of training multiple models is usually very complex compared to the simplicity of a single generative LLM framework.

2.4 Large Language Model (LLM) based Methods

Inspired by the remarkable zero-shot capabilities of LLM across various AI domains, leveraging LLM in graph problem-solving has attracted considerable attention from researchers. Existing works have already included utilizing LLM to automatically select the most suitable graph processor based on the query [57], employing LLM’s zero-shot predictions and corresponding explanations for data augmentation to obtain state-of-the-art TAPE graph feature embeddings [28], generating prompts to address graph construction problems [58], structural reasoning tasks [59], and molecular property prediction tasks [60]. Moreover, new graph problem datasets and benchmarks based on LLM have been collected and released [61]. There are three works that share similarities with our method. Guo et al. [61] attempts to complete graph tasks by describing graphs. However, it does not use natural language like our proposal. Instead, it uses complex formal languages like Brandes et al. [62] and Himsolt [63]. Wang et al. [64] and Chen et al. [65] both explore using natural language with LLM to tackle graph problems, with [64] focusing more on mathematical problems on small graphs while [65] concentrating on node classification in Text-Attributed Graphs (TAGs) [66]. Compared to Wang et al. [64] and Chen et al. [65], our designed natural language instruction prompts exhibit better regularity and scalability, applicable to both small and large graphs and not limited to specific types of graph data. In contrast, certain natural language templates in the aforementioned works are generated by LLM for specific task instructions, also enhanced with advanced prompting techniques such as Chain-of-Thoughts (CoT) [67]. Overall, the three related works only explored the basic capability of utilizing LLM for graph tasks in a zero-shot setting. Since they do not employ instruction tuning [37], their performance does not surpass GNN baselines for most of the time, only demonstrating the potential of LLM as an option for graph tasks. By contrast, our work successfully bridges this gap by conducting instruction tuning on generative LLMs with simple prompts, achieving experimental results that surpass competitive GNN baselines.

3 InstructGLM

In this section, we introduce the details of our proposed **Instruction**-finetuned **Graph Language Model**, i.e. **InstructGLM**, a framework utilizing natural language to describe both graph structure and node features to a generative large language model and further addresses graph-related problems by instruction-tuning. We start with the notation setup, followed by an introduction to the instruction prompts and their design principles, and then we explain the proposed pipeline with greater details.

3.1 Preliminary

Formally, a graph can be represented as $\mathcal{G} = (\mathcal{V}, \mathcal{A}, E, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E})$, where \mathcal{V} is the set of nodes, E is the set of edges, $\mathcal{A} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the adjacent matrix, \mathcal{N}_v is the node feature of $v \in \mathcal{V}$ and \mathcal{E}_e is the edge feature of $e \in E$. It is worth noting that the node feature and edge feature can be various modalities in diverse forms. For example, node feature can be textual information in citation networks or social networks, visual images in photography graphs, user profile in customer systems, and even video or audio signals in movie networks, while edge feature can be product reviews in user-item interaction graph of recommender systems.

3.2 Instruction Prompt Design

In order to comprehensively convey the structural information of a graph and ensure the adaptability of the created instruction prompts to various types of graphs, we have systematically designed a set of graph description prompts centered around an central node. These prompts can be differentiated based on the following three questions: **i)** What is the largest hop level of neighbor information about the central node in the prompt? **ii)** Does the prompt include node features or edge features? **iii)** For prompts with large (≥ 2) hop level neighbors about the central node, does the prompt encompass information about the intermediate nodes or paths along the corresponding connecting route?

Regarding the first question, prompts can be classified into two types: those exclusively contain 1-hop connection information, and those with a maximum of 2-hop or 3-hop connection details. Prior works have shown that utilizing up to 3-hop connectivity is sufficient for excellent performance [11–13], while information beyond 3-hop typically owns a minor impact on improvement and might even lead to negative effects [31, 68]. Therefore, the maximum level of neighbor information included in the prompts is up to three. However, benefiting from the flexibility of natural language, our designed prompts can actually accommodate structural information of any hop level. As for the latter two questions, there are two possible scenarios for each question, i.e., if or not to include the node or edge features in the prompt, and if or not to include the connecting route information in the prompt.

We then denote an instruction prompt as $\mathcal{T}(\cdot)$ such that $\mathcal{I} = \mathcal{T}(v, \mathcal{A}, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E})$ is the input sentence to LLM and v is the **central node** of this prompt with its corresponding graph structure described in natural language. For instance, the simplest form of a graph description containing at most 2-hops neighbor details is:

$$\mathcal{T}(v, \mathcal{A}) = \{v\} \text{ is connected with } \{[v_2]_{v_2 \in \mathcal{A}_2^v}\} \text{ within two hops.} \quad (1)$$

while its most detailed form which includes node features, edge features and corresponding intermediate paths should be:

$$\begin{aligned} \mathcal{T}(v, \mathcal{A}, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E}) = \{ & (v, \mathcal{N}_v)\} \text{ is connected with } \{[(v_2, \mathcal{N}_{v_2})]_{v_2 \in \mathcal{A}_2^v}\} \\ & \text{within two hops through } \{[(v_1, \mathcal{N}_{v_1})]_{v_1 \in \mathcal{A}_1^v}\} \text{ and featured} \\ & \text{paths } \{[(\mathcal{E}_{(v, v_1)}, \mathcal{E}_{(v_1, v_2)})]_{v_1 \in \mathcal{A}_1^v, v_2 \in \mathcal{A}_1^{v_1}}\}, \text{ respectively.} \end{aligned} \quad (2)$$

where \mathcal{A}_k^v represents the list of node v 's k -hop neighbor nodes. Essentially, the above prompt contains all 2-hop paths with the node and edge features such as $(v, \mathcal{N}_v) \xrightarrow{\mathcal{E}_{(v, v_1)}} (v_1, \mathcal{N}_{v_1}) \xrightarrow{\mathcal{E}_{(v_1, v_2)}} (v_2, \mathcal{N}_{v_2})$ centering at node v . All instruction prompts we designed are summarized in Appendix A.

3.3 Generative Instruction Tuning for Node Classification

In prompt engineering [69–71] or in-context learning [72], pre-trained models are usually frozen, hindering them achieving top performance in downstream tasks. Instruction Tuning [37, 38], however,

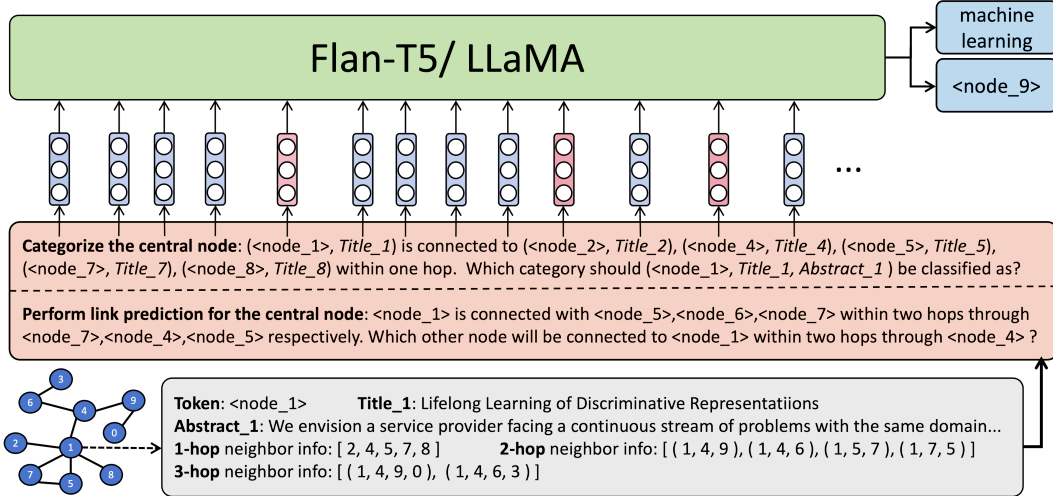


Figure 2: Illustration of InstructGLM. We use graph prompts to describe each node’s multi-hop connectivity and meta features in a scalable mini-batch manner, conveying graph structure concisely and intuitively by pure natural language for learning. Subsequently, we instruct LLM to generate responses for various graph learning tasks in a unified language modeling pipeline. We also expand LLM’s vocabulary by creating a new and unique token for every node. More specifically, we set the graph’s inherent node feature vectors (e.g. BoW, OGB) as the embedding for these new tokens (depicted as red vectors in the figure) and employ LLM’s pre-trained embedding (depicted as blue vectors in the figure) for natural language tokens.

directly conveys the requirements of downstream tasks to pre-trained models by fusing the original input data with task-specific instructional prompts under the framework of multi-prompt training. This facilitates remarkably effective fine-tuning, especially when coupled with human feedback [18]. Instruction Tuning has already become an indispensable technique for fine-tuning the most powerful large language models.

In this paper, we introduce InstructGLM as a multi-prompt instruction-tuning framework for graph learning. Specifically, we employ a generative large language model with an encoder-decoder or decoder-only architecture as the backbone, then fuse all of our designed instruction prompts, which are spanning at different hop levels with diverse structural information, together as input to LLM, enabling mutual enhancement among the instructions. By exclusively using natural language to depict graph structures, we succinctly present the graph’s geometry to the LLM and provide a pure NLP interface for all graph-related tasks, make them solvable through a unified pipeline in generative manner. Worth noting that we concentrate on solving node classification task in this study. We train InstructGLM to strictly generate the category label in natural language, and the prevalent Negative Log-Likelihood (i.e. NLL) Loss in language modeling are selected as our objective function.

Formally, given graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, E, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E})$ and a specific instruction prompt $\mathcal{T} \in \{\mathcal{T}(\cdot)\}$, we denote \mathbf{x} and \mathbf{y} as LLM’s input and target sentence, respectively. Then our pipeline can be formed as:

$$P_\theta(\mathbf{y}_j | \mathbf{x}, \mathbf{y}_{<j}) = \text{LLM}_\theta(\mathbf{x}, \mathbf{y}_{<j}), \quad \mathbf{x} = \text{Concatenate}(\mathcal{P}; \mathcal{I}; \mathcal{Q}) \quad (3)$$

$$\mathcal{L}_\theta = - \sum_{j=1}^{|\mathbf{y}|} \log P_\theta(\mathbf{y}_j | \mathbf{x}, \mathbf{y}_{<j}) \quad (4)$$

where \mathcal{L} denotes the NLL loss, $\mathcal{I} = \mathcal{T}(v, \mathcal{A}, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E})$ is the graph structure description centering at node $v \in \mathcal{V}$, \mathcal{P} and \mathcal{Q} are the task-specific instruction prefix and query. Specifically, for node classification, we design \mathcal{P} and \mathcal{Q} for node classification as follows: $\mathcal{P} = \text{‘Classify the central node into one of the following categories: } [\langle \text{All category} \rangle]. \text{ Pay attention to the multi-hop link relationships between the nodes.} \text{’}$ and $\mathcal{Q} = \text{‘Which category should } \{v\} \text{ be classified as?’}$. More details of the pipeline are depicted in Figure 2.

Our InstructGLM actually shares essential similarities in mechanism with various GNNs, and thus covering their advantages. First, we mix prompts with diverse hop-level information together during training, which is akin to MixHop [35] in performing graph convolutions on subgraphs extracted at different hop levels. Second, Jumping Knowledge [36] combines outcomes from different convolution layers via jump connections, which is aligned with our prompts featuring intermediate information and high-hop-level neighbors. Additionally, due to LLM’s input length limit, similar to GraphSAGE [13], we conduct neighbor sampling for the central node when filling the prompts to form a mini-batch training. This operation also resembles graph regularization techniques like DropEdge [46] for preventing over-smoothing [73].

Furthermore, compared to GNNs, our InstructGLM exhibits stronger expressive capabilities. In our method, even a single graph description that contains intermediate paths and k -hop neighbor information is equivalent to a k -layer GNN in expressiveness. Therefore, InstructGLM can readily accommodate the inductive bias of graph tasks without any alterations on LLM’s architecture and pipeline. For instance, since our inputs are centralized graph descriptions that directly exhibit the corresponding multi-hop neighbors, self-attention [1] applied on such inputs can be seen as an advanced weighted average aggregation mechanism of GATs [12, 74], facilitating InstructGLM to effectively grasp different neighbors’ varying importance to the central node.

3.4 Auxiliary Self-Supervised Link Prediction

Both SuperGAT [75] and DiffPool [76] introduce auxiliary link prediction task, thus successfully obtain better node representations and performance for node or graph classification, demonstrating that model’s comprehension of graph structure can be significantly enhanced by such an auxiliary task. Inspired by them, also to remove the restriction that our instruction prompts can only treat labeled training nodes as central nodes in single-task semi-supervised learning, we introduce self-supervised link prediction as a foundational auxiliary task for InstructGLM. Given arbitrary hop level, for every node in the graph, we can randomly select a neighbor or non-neighbor at this hop level as its candidate. Then we instruct our model to either discriminate whether there is a connection at this hop level between the central node and the candidate node (discriminative prompt) or directly generate the correct neighbor in a generative manner (generative prompt).

More formally, given graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, E, \{\mathcal{N}_v\}_{v \in \mathcal{V}}, \{\mathcal{E}_e\}_{e \in E})$, the pipeline of link prediction aligns exactly with the aforementioned equations 3 and 4. The only distinction lies in the newly designed task-specific prefix and two different query templates for it. Specifically, we design \mathcal{P} and \mathcal{Q} for link prediction as follows: $\mathcal{P} = \text{‘Perform link prediction for the central node. Pay attention to the multi-hop link relationships between the nodes.’}$, $\mathcal{Q}_{generative} = \text{‘Which other node will be connected to } \{v\} \text{ within } \{h\} \text{ hop?’}$ and $\mathcal{Q}_{discriminative} = \text{‘Will } \{\tilde{v}\} \text{ be connected to } \{v\} \text{ within } \{h\} \text{ hop?’}$, where v is the central node, \tilde{v} is the candidate node and h is the specified hop level.

Consequently, we extend InstructGLM into a multi-task, multi-prompt instruction-tuning framework. Irrespective of the graph type and graph-related task that InstructGLM ultimately aims at, the inclusion of auxiliary self-supervised link prediction enables every node in the graph to act as the central node in multiple instruction prompts during training. Thus, it not only serves as data augmentation but also encourages the LLM to understand the graph’s global connectivity pattern, providing InstructGLM with promising potential for further performance improvement on the primary task.

4 Experiments

4.1 Experimental Setup

In this paper, we primarily utilize InstructGLM for node classification, also conduct self-supervised link prediction as an auxiliary task. Specifically, we select the following three popular graphs: ogbn-arxiv [66], Cora, and PubMed [77], in which every node represents an academic paper on a specific topic, with its title and abstract included in raw text format, and there will be an edge between the corresponding two nodes if there exists a citation between two papers. The graph of ogbn-arxiv is relatively larger while the graphs of Cora and PubMed are smaller. All of our experiments employ the default numerical node feature embedding provided by the dataset to extend the LLM’s vocabulary by adding node-wise newly constructed tokens. It is worth noting that these datasets used different techniques to generate the default node feature embeddings and we used their default embeddings without modification. Detailed dataset-specific information is summarized in Table 1.

Table 1: Dataset Statistics

Dataset	#Node	#Edge	#Class	Default Feature	#Features
ogbn-arxiv	169,343	1,166,243	40	Skip-gram / GIANT	128 / 768
Cora	2,708	5,429	7	Bag of Words	1433
PubMed	19,717	44,338	3	TF-IDF	500

We employ a multi-prompt instruction-tuning framework for all of our experiments and report test accuracy as our metrics. For ogbn-arxiv dataset, we adopt exactly the same dataset splits as in the OGB open benchmark [66], which is 54%/18%/28%. For Cora and PubMed datasets, we use the version that contains raw text information proposed in He et al. [28] and employ a 60%/20%/20% train/val/test splits for our experiments. To investigate InstructGLM’s performance under low-label-ratio training setting, following Yang et al. [77], we conduct further experiments on the PubMed dataset with the fixed 20 labeled training nodes per class at a 0.3% label ratio.

4.2 Main Results

Our results achieve state-of-the-art performance, surpassing all single-model graph learners across all three datasets, including both representative GNN models and graph Transformer models, which demonstrates the promising trend for large language models to serve as the foundation model for graph learning. More detailed results and analyses are presented in the following.

4.2.1 ogbn-arxiv

For the ogbn-arxiv dataset, we select a series of top-ranked GNNs from the OGB Leaderboard¹, including DRGAT, RevGAT, AGDN, etc., as the baselines. And several most powerful Transformer-based single-model graph learners like Graphormer and E2EG on the OGB benchmark are also considered as compared methods against our proposed InstructGLM.

Table 2: Results on ogbn-arxiv

Method		OGB	GIANT
MLP		55.50 ± 0.23	73.06 ± 0.11
GAMLP	(Zhang et al. [78])	56.53 ± 0.16	73.35 ± 0.08
GraphSAGE	(Hamilton et al. [13])	71.19 ± 0.21	74.35 ± 0.14
GCN	(Kipf and Welling [11])	71.74 ± 0.29	73.29 ± 0.01
DeeperGCN	(Li et al. [79])	71.92 ± 0.16	–
ALT-OPT	(Han et al. [80])	72.76 ± 0.00	–
GTAN	(Wu and Wang [81])	72.97 ± 0.17	–
UniMP	(Shi et al. [82])	73.11 ± 0.20	–
LEGNN	(Yu et al. [83])	73.37 ± 0.07	–
GAT	(Veličković et al. [12])	73.66 ± 0.11	74.15 ± 0.05
AGDN	(Sun et al. [84])	73.75 ± 0.21	76.02 ± 0.16
RvGAT	(Li et al. [74])	74.02 ± 0.18	75.90 ± 0.19
DRGAT	(Zhang et al. [85])	74.16 ± 0.07	<u>76.11 ± 0.09</u>
NodeFormer	(Wu et al. [86])	59.90 ± 0.42	–
CoarFormer	(Kuang et al. [87])	71.66 ± 0.24	–
SGFormer	(Wu et al. [88])	72.63 ± 0.13	–
Graphormer	(Ying et al. [23])	72.81 ± 0.23	–
E2EG	(Dinh et al. [89])	73.62 ± 0.14	–
InstructGLM-Flan-T5-base (ours)		73.51 ± 0.16	74.45 ± 0.11
InstructGLM-Flan-T5-large (ours)		<u>74.67 ± 0.08</u>	74.80 ± 0.18
InstructGLM-Llama-v1-7b (ours)		75.70 ± 0.12	76.42 ± 0.09

We instruction-finetune Flan-T5 [38] and Llama-v1-7b [17] (LoRA) [90] as the backbone for our InstructGLM. The experimental results in Table 2 demonstrate that InstructGLM outperforms all the

¹https://ogb.stanford.edu/docs/leader_nodeprop/

GNNs and Transformer-based methods. Particularly, when using Llama-v1-7b as the backbone on the OGB feature, our InstructGLM attains a **1.54%** improvement over the best GNN method and a **2.08%** improvement over the best Transformer-based method. Meanwhile, we also obtain new **SoTA** performance on the GIANT [54] feature.

4.2.2 Cora & PubMed

In terms of the compared methods for Cora and PubMed datasets, we select those top-ranked GNNs from the two corresponding benchmarks^{2,3}, including Snowball, MixHop, RevGAT, FAGCN, etc., as the baselines. Besides, the three most powerful Transformer-based single-model graph learners on these 2 benchmarks, i.e., CoarFormer, Graphormer, and GT, are also considered as compared methods against our proposed InstructGLM.

Table 3: Results on Cora and PubMed

Cora	Acc	PubMed	Acc
MixHop (Abu-El-Haija et al. [35])	75.65 ± 1.31	GAT (Veličković et al. [12])	83.28 ± 0.12
GAT (Veličković et al. [12])	76.70 ± 0.42	SGC-v2 (Wu et al. [91])	85.36 ± 0.52
Geom-GCN (Pei et al. [92])	85.27 ± 1.48	GraphSAGE (Hamilton et al. [13])	86.85 ± 0.11
SGC-v2 (Wu et al. [91])	85.48 ± 1.48	BernNet (He et al. [93])	88.48 ± 0.41
GraphSAGE (Hamilton et al. [13])	86.58 ± 0.26	RevGAT (Li et al. [74])	88.50 ± 0.05
GCN (Kipf and Welling [11])	87.78 ± 0.96	GCN (Kipf and Welling [11])	88.90 ± 0.32
BernNet (He et al. [93])	88.52 ± 0.95	GCNII (Chen et al. [94])	89.80 ± 0.30
FAGCN (Bo et al. [95])	88.85 ± 1.36	FAGCN (Bo et al. [95])	89.98 ± 0.54
GCNII (Chen et al. [94])	88.93 ± 1.37	MixHop (Abu-El-Haija et al. [35])	90.04 ± 1.41
RevGAT (Li et al. [74])	89.11 ± 0.00	Geom-GCN (Pei et al. [92])	90.05 ± 0.14
ACM-Snowball-3 (Luan et al. [96])	89.59 ± 1.58	ACM-GCN+ (Luan et al. [96])	90.96 ± 0.62
ACM-GCN+ (Luan et al. [96])	89.75 ± 1.16	ACM-Snowball-3 (Luan et al. [96])	91.44 ± 0.59
Graphormer (Ying et al. [23])	80.41 ± 0.30	GT (Dwivedi and Bresson [33])	88.75 ± 0.16
GT (Dwivedi and Bresson [33])	86.42 ± 0.82	Graphormer (Ying et al. [23])	88.24 ± 1.50
CoarFormer (Kuang et al. [87])	88.69 ± 0.82	CoarFormer (Kuang et al. [87])	89.75 ± 0.31
InstructGLM-Llama-v1-7b (ours)	87.08 ± 0.32	InstructGLM-Llama-v1-7b (ours)	93.84 ± 0.25
InstructGLM-Flan-T5-base (ours)	90.77 ± 0.52	InstructGLM-Flan-T5-base (ours)	94.45 ± 0.12
InstructGLM-Flan-T5-large (ours)	88.93 ± 1.06	InstructGLM-Flan-T5-large (ours)	94.62 ± 0.13

We instruction-finetune Flan-T5 and Llama-v1 (LoRA) as the backbone for our InstructGLM. The experimental results in Table 3 show that our InstructGLM outperforms all the GNNs and Transformer-based methods. Specifically, InstructGLM achieves a **1.02%** improvement over the best GNN method and a **2.08%** improvement over the best Transformer-based method on the Cora dataset, while also achieves a **3.18%** improvement over the best GNN method and a **4.87%** improvement over the best Transformer-based method on the PubMed dataset.

4.3 Ablation Study

In our experiments, two crucial operations that contributes to the remarkable performance of InstructGLM in node classification are multi-prompt instruction-tuning, which provides multi-hop graph structure information to the LLM, and the utilization of self-supervised link prediction as an auxiliary task. To validate the impact of the two key components on model performance, we conduct ablation experiments on all three datasets, the results are shown in Table 4.

Table 4: Ablation Study Results

Hop Info	Link Prediction	ogbn-arxiv	Cora	PubMed
		Llama-v1-7b	Flan-T5-base	Flan-T5-base
Multi-hop	w/	75.70%	90.77%	94.45%
Multi-hop	w/o	75.37%	87.27%	94.35%
1-hop	w/o	75.25%	86.90%	94.30%
Structure-Free-Tuning	w/o	74.97%	75.65%	94.22%

²<https://paperswithcode.com/sota/node-classification-on-cora-60-20-20-random>

³<https://paperswithcode.com/sota/node-classification-on-pubmed-60-20-20-random>

Regarding the *Hop Info* column, *Structure-Free-Tuning* indicates that we do not consider the graph’s structure, which means directly fine-tuning the model on titles and abstracts of the nodes. While *1-hop* and *Multi-hop* mean that we utilize prompts that merely include information from 1-hop neighbors and prompts that include information from neighbors with higher hop levels, respectively. The experimental results show that incorporating multi-hop information and including link prediction task can both enhance the model’s performance in node classification tasks.

4.4 Instruction Tuning at Low Label Ratio

In previous experiments, our data splits all ensured a relatively high ratio of labeled training nodes. To further investigate the scalability and robustness of our InstructGLM, we conduct experiments on the PubMed dataset using its another widely-used splits with extremely low label ratio. Specifically, we have only 60 training nodes available in this setting thus the label ratio is **0.3%**.

We consider top-ranked GNNs from the corresponding leaderboard⁴, including SAIL, ALT-OPT, GRAND, etc., as the baselines to be compared with our InstructGLM. We also take the three most outstanding Transformer-based graph learners under this dataset setting into account. We then instruction-finetune Flan-T5 and Llama as the backbone for our InstructGLM. The experimental results in Table 5 demonstrate that InstructGLM outperforms all the GNNs methods with an improvement of **5.8%** against the best GNN baseline, while also surpassing the best Transformer-based model by **9.3%**, successfully achieve new state-of-the-art performance on the leaderboard.

Table 5: Results on PubMed with 60 training nodes

Method	Acc
GraphSAGE (Hamilton et al. [13])	76.8 ± 0.9
GAT (Veličković et al. [12])	79.0 ± 1.4
Snowball (Luan et al. [97])	79.2 ± 0.3
GCN (Kipf and Welling [11])	80.4 ± 0.4
SuperGAT (Kim and Oh [75])	81.7 ± 0.5
ALT-OPT (Han et al. [80])	82.5 ± 1.7
GRAND (Feng et al. [98])	82.7 ± 0.6
SAIL (Yu et al. [99])	83.8 ± 0.1
ANS-GT (Zhang et al. [100])	79.6 ± 1.0
NodeFormer (Wu et al. [86])	79.9 ± 1.0
SGFormer (Wu et al. [88])	80.3 ± 0.6
InstructGLM-Llama-v1-7b (ours)	85.1 ± 0.6
InstructGLM-Flan-T5-base (ours)	88.2 ± 0.3
InstructGLM-Flan-T5-large (ours)	89.6 ± 0.4

5 Future Work

In this paper, we conduct extensive experiments on Text-Attributed Graphs (TAG) to showcase the powerful capabilities of our proposed InstructGLM in solving graph machine learning problems. Our instruction prompts designed to describe graph structures in natural language demonstrate high generality and scalability, making them applicable to almost all types of graphs. Potential valuable future work can be explored along three dimensions:

- For TAGs, our experiments only used the default OGB-feature embeddings. Future work can consider using more advanced TAG-related embedding features such as LLM-based features like TAPE [28]. Additionally, leveraging LLM for Chain-of-Thought, structure information summary, and other data augmentation techniques to generate more powerful instruction prompts will be a promising research direction for graph language models.
- InstructGLM can be integrated into frameworks like GAN, GLEM for multi-model iterative training, or utilize off-the-shelf GNNs for knowledge distillation. Also, classic graph machine learning techniques like label reuse, Self-KD, Correct & Smooth can further enhance the model’s performance.
- Benefiting from the powerful expressive ability of natural language and the highly scalable design of our instruction prompts, InstructGLM can be easily extended within a unified generative language modeling framework to various kinds of graphs, addressing a wide range of graph learning problems. For instance, our designed instruction prompts can be directly further used for link prediction and inductive node classification tasks. And only with slight modifications to our prompts, tasks such as graph classification, intermediate node/path prediction and even relation-based question answering tasks in knowledge graphs with rich edge features can be effectively deployed.

⁴<https://paperswithcode.com/sota/node-classification-on-pubmed-with-public>

6 Conclusion

To the best of our knowledge, this paper is the first one that purely represents graph structure via natural language description then further perform instruction-tuning on generative large language models to effectively solve graph learning problems, demonstrating the huge potential of LLMs to replace GNNs as the new foundational model for graph machine learning. Leveraging the powerful expressive ability of natural language, our proposed InstructGLM outperforms all competitive GNNs and Transformer-based graph learners on the ogbn-arxiv, Cora, and PubMed datasets under the default feature embeddings setting, showcasing its strong capability in solving graph problems. Moreover, benefiting from our highly scalable instruction prompts and the unified generative pipeline that is applicable to multi-modality data, InstructGLM can be readily extended to promising future work as discussed above. Overall, our InstructGLM provides a powerful natural language processing interface for graph machine learning problems, with Transformer-based generative large model and natural language as the driving force, it further contributes to the trend of unifying foundational model architecture and pipeline across multi-modality for the AGI pursuit in the future.

References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1, 3, 7
- [2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [3] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 1
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1
- [5] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1
- [6] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. 1
- [7] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1
- [8] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 1
- [9] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Super-glue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. 1
- [10] I-Fan Chen, Brian King, and Jasha Droppo. Investigation of training label error impact on rnn-t. *arXiv preprint arXiv:2112.00350*, 2021. 1
- [11] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 1, 3, 5, 8, 9, 10
- [12] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 3, 7, 8, 9, 10
- [13] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017. 1, 3, 5, 7, 8, 9, 10
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1

- [15] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 1
- [16] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [17] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 8
- [18] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022. 1, 6
- [19] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023. 1
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [21] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 1
- [22] Shikhar Singh, Nuan Wen, Yu Hou, Pegah Alipoormolabashi, Te-Lin Wu, Xuezhe Ma, and Nanyun Peng. Com2sense: A commonsense reasoning benchmark with complementary sentences. *arXiv preprint arXiv:2106.00969*, 2021. 1
- [23] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34:28877–28888, 2021. 1, 4, 8, 9
- [24] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems*, pages 299–315, 2022. 1, 2
- [25] Norman Di Palo, Arunkumar Byravan, Leonard Hasenclever, Markus Wulfmeier, Nicolas Heess, and Martin Riedmiller. Towards a unified agent with foundation models. In *Workshop on Reincarnating Reinforcement Learning at ICLR 2023*, 2023. 1, 2
- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2
- [27] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-transformer: A unified framework for multimodal learning. *arXiv preprint arXiv:2307.10802*, 2023. 1
- [28] Xiaoxin He, Xavier Bresson, Thomas Laurent, and Bryan Hooi. Explanations as features: Llm-based features for text-attributed graphs. *arXiv preprint arXiv:2305.19523*, 2023. 2, 4, 8, 10
- [29] Costas Mavromatis, Vassilis N Ioannidis, Shen Wang, Da Zheng, Soji Adeshina, Jun Ma, Han Zhao, Christos Faloutsos, and George Karypis. Train your own gnn teacher: Graph-aware distillation on textual graphs. *arXiv preprint arXiv:2304.10668*, 2023. 4
- [30] Jianan Zhao, Meng Qu, Chaozhuo Li, Hao Yan, Qian Liu, Rui Li, Xing Xie, and Jian Tang. Learning on large-scale text-attributed graphs via variational inference. *ICLR*, 2023. 2, 4

- [31] Chen Cai and Yusu Wang. A note on over-smoothing for graph neural networks. *arXiv preprint arXiv:2006.13318*, 2020. 2, 3, 5
- [32] Jinwoo Kim, Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. Pure transformers are powerful graph learners. *Advances in Neural Information Processing Systems*, 35:14582–14595, 2022. 2, 4
- [33] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020. 2, 9
- [34] Dai Quoc Nguyen, Tu Dinh Nguyen, and Dinh Phung. Universal graph transformer self-attention networks. In *Companion Proceedings of the Web Conference 2022*, pages 193–196, 2022. 2, 4
- [35] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*, pages 21–29. PMLR, 2019. 3, 7, 9
- [36] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *International conference on machine learning*, pages 5453–5462. PMLR, 2018. 3, 7
- [37] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021. 3, 4, 5
- [38] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. 3, 5, 8
- [39] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018. 3
- [40] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI open*, 1:57–81, 2020. 3
- [41] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020. 3
- [42] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018. 3
- [43] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124, 2017. 3
- [44] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016. 3
- [45] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. Hypergc: A new method for training graph convolutional networks on hypergraphs. *Advances in neural information processing systems*, 32, 2019. 3
- [46] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. Droppedge: Towards deep graph convolutional networks on node classification. *arXiv preprint arXiv:1907.10903*, 2019. 3, 7
- [47] Yangkun Wang, Jiarui Jin, Weinan Zhang, Yong Yu, Zheng Zhang, and David Wipf. Bag of tricks for node classification with graph neural networks. *arXiv preprint arXiv:2103.13355*, 2021. 3
- [48] Luis Müller, Mikhail Galkin, Christopher Morris, and Ladislav Rampásek. Attending to graph transformers. *arXiv preprint arXiv:2302.04181*, 2023. 4
- [49] Wonpyo Park, Woonggi Chang, Donggeon Lee, Juntae Kim, and Seung-won Hwang. Grpe: Relative positional encoding for graph transformer. *arXiv preprint arXiv:2201.12787*, 2022. 4

- [50] Jiayou Zhang, Zhirui Wang, Shizhuo Zhang, Megh Manoj Bhalerao, Yucong Liu, Dawei Zhu, and Sheng Wang. Graphprompt: Biomedical entity normalization using graph-based prompt templates. *arXiv preprint arXiv:2112.03002*, 2021. 4
- [51] Robert R Sokal and Theodore J Crovello. The biological species concept: a critical evaluation. *The American Naturalist*, 104(936):127–153, 1970. 4
- [52] Qinyong Wang, Zhenxiang Gao, and Rong Xu. Exploring the in-context learning ability of large language model for biomedical concept linking. *arXiv preprint arXiv:2307.01137*, 2023. 4
- [53] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 4
- [54] Eli Chien, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Jiong Zhang, Olgica Milenkovic, and Inderjit S Dhillon. Node feature extraction by self-supervised multi-scale neighborhood prediction. *arXiv preprint arXiv:2111.00064*, 2021. 4, 9
- [55] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33:12559–12571, 2020. 4
- [56] Chaoqi Yang, Ruijie Wang, Shuochao Yao, Shengzhong Liu, and Tarek Abdelzaher. Revisiting over-smoothing in deep gcns. *arXiv preprint arXiv:2003.13663*, 2020. 4
- [57] Jiawei Zhang. Graph-toolformer: To empower llms with graph reasoning ability via prompt augmented by chatgpt. *arXiv preprint arXiv:2304.11116*, 2023. 4
- [58] Jiuzhou Han, Nigel Collier, Wray Buntine, and Ehsan Shareghi. Pive: Prompting with iterative verification improving graph-based generative capability of llms. *arXiv preprint arXiv:2305.12392*, 2023. 4
- [59] Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Wayne Xin Zhao, and Ji-Rong Wen. Structgpt: A general framework for large language model to reason over structured data. *arXiv preprint arXiv:2305.09645*, 2023. 4
- [60] Chen Qian, Huayi Tang, Zhirui Yang, Hong Liang, and Yong Liu. Can large language models empower molecular property prediction? *arXiv preprint arXiv:2307.07443*, 2023. 4
- [61] Jiayan Guo, Lun Du, and Hengyu Liu. Gpt4graph: Can large language models understand graph structured data? an empirical evaluation and benchmarking. *arXiv preprint arXiv:2305.15066*, 2023. 4
- [62] Ulrik Brandes, Markus Eiglsperger, Jürgen Lerner, and Christian Pich. Graph markup language (graphml), 2013. 4
- [63] Michael Himsolt. Gml: A portable graph file format. Technical report, Technical report, Universitat Passau, 1997. 4
- [64] Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. Can language models solve graph problems in natural language? *arXiv preprint arXiv:2305.10037*, 2023. 4
- [65] Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Hongzhi Wen, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, Wenqi Fan, Hui Liu, et al. Exploring the potential of large language models (llms) in learning on graphs. *arXiv preprint arXiv:2307.03393*, 2023. 4
- [66] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020. 4, 7, 8
- [67] 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. 4
- [68] Wentao Zhang, Zeang Sheng, Yuezihan Jiang, Yikuan Xia, Jun Gao, Zhi Yang, and Bin Cui. Evaluating deep graph neural networks. *arXiv preprint arXiv:2108.00955*, 2021. 5
- [69] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 5

- [70] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- [71] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Auto-prompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020. 5
- [72] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*, 2022. 5
- [73] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 3438–3445, 2020. 7
- [74] Guohao Li, Matthias Müller, Bernard Ghanem, and Vladlen Koltun. Training graph neural networks with 1000 layers. In *International conference on machine learning*, pages 6437–6449. PMLR, 2021. 7, 8, 9
- [75] Dongkwan Kim and Alice Oh. How to find your friendly neighborhood: Graph attention design with self-supervision. *arXiv preprint arXiv:2204.04879*, 2022. 7, 10
- [76] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. *Advances in neural information processing systems*, 31, 2018. 7
- [77] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pages 40–48. PMLR, 2016. 7, 8
- [78] Wentao Zhang, Ziqi Yin, Zeang Sheng, Yang Li, Wen Ouyang, Xiaosen Li, Yangyu Tao, Zhi Yang, and Bin Cui. Graph attention multi-layer perceptron. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4560–4570, 2022. 8
- [79] Guohao Li, Chenxin Xiong, Ali Thabet, and Bernard Ghanem. Deepergcn: All you need to train deeper gcns. *arXiv preprint arXiv:2006.07739*, 2020. 8
- [80] Haoyu Han, Xiaorui Liu, Haitao Mao, MohamadAli Torkamani, Feng Shi, Victor Lee, and Jiliang Tang. Alternately optimized graph neural networks. In *International Conference on Machine Learning*, pages 12411–12429. PMLR, 2023. 8, 10
- [81] Nan Wu and Chaofan Wang. Gtnet: A tree-based deep graph learning architecture. *arXiv preprint arXiv:2204.12802*, 2022. 8
- [82] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*, 2020. 8
- [83] Le Yu, Leilei Sun, Bowen Du, Tongyu Zhu, and Weifeng Lv. Label-enhanced graph neural network for semi-supervised node classification. *IEEE Transactions on Knowledge and Data Engineering*, 2022. 8
- [84] Chuxiong Sun, Jie Hu, Hongming Gu, Jinpeng Chen, and Mingchuan Yang. Adaptive graph diffusion networks. *arXiv preprint arXiv:2012.15024*, 2020. 8
- [85] Lei Zhang, Xiaodong Yan, Jianshan He, Ruopeng Li, and Wei Chu. Drgcn: Dynamic evolving initial residual for deep graph convolutional networks. *arXiv preprint arXiv:2302.05083*, 2023. 8
- [86] Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. Nodeformer: A scalable graph structure learning transformer for node classification. *Advances in Neural Information Processing Systems*, 35:27387–27401, 2022. 8, 10
- [87] Weirui Kuang, WANG Zhen, Yaliang Li, Zhewei Wei, and Bolin Ding. Coarformer: Transformer for large graph via graph coarsening. 2021. 8, 9
- [88] Qitian Wu, Wentao Zhao, Chenxiao Yang, Hengrui Zhang, Fan Nie, Haitian Jiang, Yatao Bian, and Junchi Yan. Simplifying and empowering transformers for large-graph representations. *arXiv preprint arXiv:2306.10759*, 2023. 8, 10

- [89] Tu Anh Dinh, Jeroen den Boef, Joran Cornelisse, and Paul Groth. E2eg: End-to-end node classification using graph topology and text-based node attributes. *arXiv preprint arXiv:2208.04609*, 2022. 8
- [90] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 8
- [91] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019. 9
- [92] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020. 9
- [93] Mingguo He, Zhewei Wei, Hongteng Xu, et al. Bernnet: Learning arbitrary graph spectral filters via bernstein approximation. *Advances in Neural Information Processing Systems*, 34: 14239–14251, 2021. 9
- [94] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *International conference on machine learning*, pages 1725–1735. PMLR, 2020. 9
- [95] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3950–3957, 2021. 9
- [96] Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. Revisiting heterophily for graph neural networks. *Advances in neural information processing systems*, 35:1362–1375, 2022. 9
- [97] Sitao Luan, Mingde Zhao, Xiao-Wen Chang, and Doina Precup. Break the ceiling: Stronger multi-scale deep graph convolutional networks. *Advances in neural information processing systems*, 32, 2019. 10
- [98] Wenzheng Feng, Jie Zhang, Yuxiao Dong, Yu Han, Huanbo Luan, Qian Xu, Qiang Yang, Evgeny Kharlamov, and Jie Tang. Graph random neural networks for semi-supervised learning on graphs. *Advances in neural information processing systems*, 33:22092–22103, 2020. 10
- [99] Lu Yu, Shichao Pei, Lizhong Ding, Jun Zhou, Longfei Li, Chuxu Zhang, and Xiangliang Zhang. Sail: Self-augmented graph contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8927–8935, 2022. 10
- [100] Zaixi Zhang, Qi Liu, Qingyong Hu, and Chee-Kong Lee. Hierarchical graph transformer with adaptive node sampling. *Advances in Neural Information Processing Systems*, 35:21171–21183, 2022. 10

A Appendix

A.1 Instruction Prompts

In this appendix, we present all our designed instruction prompts. It is worth noting that we follow the following conventions when numbering the prompts:

- The length of each prompt number is 4.
- The first digit represents the task index, where 1 represents the node classification task and 2 represents the link prediction task.
- The second digit represents whether node features or edge features (such as text information) other than numerical feature embedding are used in the prompt. 1 means not used and 2 means used.
- The third digit represents the maximum hop order corresponding to the structural information considered in this prompt. 1 represents only the 1-hop neighbors are included, while 2 and 3 represent the structural information including 2-hop and 3-hop neighbors, respectively.
- The fourth digit represents whether the intermediate node information (i.e. the path) in the high-order connection is considered in this prompt. If the digit is even, it means that the intermediate node is considered, while an odd digit indicates otherwise.
- Specially, in node classification task, we designed a graph-structure-free prompt and numbered it as 1-0-0-0.

A.4.1 Node Classification

Task-specific prefix:

Classify the paper according to its topic into one of the following categories: `{{All Categories}}`. \n Node represents academic paper with a specific topic, link represents a citation between the two papers. Pay attention to the multi-hop link relationship between the nodes.

Prompt ID: 1-1-1-1

Input template:

`{{central node}}` is connected with `{{1-hop neighbor list}}` within one hop. Which category should `{{central node}}` be classified as?

Target template: `{{category}}`

Prompt ID: 1-1-2-1

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops. Which category should `{{central node}}` be classified as?

Target template: `{{category}}`

Prompt ID: 1-1-2-2

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops through `{{the corresponding 1-hop intermediate node list}}`, respectively. Which category should `{{central node}}` be classified as?

Target template: `{{category}}`

Prompt ID: 1-1-3-1

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops. Which category should `{{central node}}` be classified as?

Target template: `{{category}}`

Prompt ID: 1-1-3-2

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops through `{{the corresponding 2-hop intermediate path list}}`, respectively. Which category should `{{central node}}` be classified as?

Target template: `{{category}}`

Prompt ID: 1-2-1-1

Input template:

(`{{central node}}`, `{{text feature}}`) is connected with `{{1-hop neighbor list attached with text feature}}` within one hop. Which category should (`{{central node}}`, `{{text feature}}`) be classified as?

Target template: `{{category}}`

Prompt ID: 1-2-2-1

Input template:

{{central node}}, {{text feature}} is connected with {{2-hop neighbor list attached with text feature}} within two hops. Which category should {{central node}}, {{text feature}} be classified as?

Target template: {{category}}

Prompt ID: 1-2-2-2

Input template:

{{central node}}, {{text feature}} is connected with {{2-hop neighbor list attached with text feature}} within two hops through {{the corresponding 1-hop intermediate node list attached with text feature}}, respectively. Which category should {{central node}}, {{text feature}} be classified as?

Target template: {{category}}

Prompt ID: 1-2-3-1

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops. Which category should {{central node}}, {{text feature}} be classified as?

Target template: {{category}}

Prompt ID: 1-2-3-2

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops through {{the corresponding 2-hop intermediate path list attached with text feature}}, respectively. Which category should {{central node}}, {{text feature}} be classified as?

Target template: {{category}}

Prompt ID: 1-0-0-0

Input template:

{{central node}} is featured with its {{text feature}}. Which category should {{central node}} be classified as?

Target template: {{category}}

A.4.2 Link Prediction

Task-specific prefix:

Perform Link Prediction for the central node:\n Node represents academic paper with a specific topic, link represents a citation between the two papers. Pay attention to the multi-hop link relationship between the nodes.

Prompt ID: 2-1-1-1

Input template:

{{central node}} is connected with {{1-hop neighbor list}} within one hop. Will {{candidate node}} be connected with {{central node}} within one hop?

Target template: {{yes/no}}

Prompt ID: 2-1-1-2

Input template:

`{{central node}}` is connected with `{{1-hop neighbor list}}` within one hop. Which other node will be connected to `{{central node}}` within one hop?

Target template: `{{node_id}}`

Prompt ID: 2-1-2-1

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops. Will `{{candidate node}}` be connected to `{{central node}}` within two hops?

Target template: `{{yes/no}}`

Prompt ID: 2-1-2-2

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops through `{{the corresponding 1-hop intermediate node list}}`, respectively. Will `{{candidate node}}` be connected to `{{central node}}` within two hops through `{{the specified 1-hop intermediate node}}`?

Target template: `{{yes/no}}`

Prompt ID: 2-1-2-3

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops. Which other node will be connected to `{{central node}}` within two hops?

Target template: `{{node_id}}`

Prompt ID: 2-1-2-4

Input template:

`{{central node}}` is connected with `{{2-hop neighbor list}}` within two hops through `{{the corresponding 1-hop intermediate node list}}`, respectively. Which other node will be connected to `{{central node}}` within two hops through `{{the specified 1-hop intermediate node}}`?

Target template: `{{node_id}}`

Prompt ID: 2-1-3-1

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops. Will `{{candidate node}}` be connected with `{{central node}}` within three hops?

Target template: `{{yes/no}}`

Prompt ID: 2-1-3-2

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops through `{{the corresponding 2-hop intermediate path list}}`, respectively. Will `{{candidate node}}` be connected to `{{central node}}` within three hops through `{{the specified 2-hop intermediate path}}`?

Target template: `{{yes/no}}`

Prompt ID: 2-1-3-3

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops. Which other node will be connected to `{{central node}}` within three hops?

Target template: `{{node_id}}`

Prompt ID: 2-1-3-4

Input template:

`{{central node}}` is connected with `{{3-hop neighbor list}}` within three hops through `{{the corresponding 2-hop intermediate path list}}`, respectively. Which other node will be connected to `{{central node}}` within three hops through `{{the specified 2-hop intermediate path}}`?

Target template: `{{node_id}}`

Prompt ID: 2-2-1-1

Input template:

`({{central node}},{{text feature}})` is connected with `{{1-hop neighbor list attached with text feature}}` within one hop. Will `({{candidate node}},{{candidate text feature}})` be connected to `({{central node}},{{text feature}})` within one hop?

Target template: `{{yes/no}}`

Prompt ID: 2-2-1-2

Input template:

`({{central node}},{{text feature}})` is connected with `{{1-hop neighbor list attached with text feature}}` within one hop. Which other node will be connected to `({{central node}},{{text feature}})` within one hop?

Target template: `{{node_id}}`

Prompt ID: 2-2-2-1

Input template:

`({{central node}},{{text feature}})` is connected with `{{2-hop neighbor list attached with text feature}}` within two hops. Will `({{candidate node}},{{candidate text feature}})` be connected to `({{central node}},{{text feature}})` within two hops?

Target template: `{{yes/no}}`

Prompt ID: 2-2-2-2

Input template:

`({{central node}},{{text feature}})` is connected with `{{2-hop neighbor list attached with text feature}}` within two hops through `{{the corresponding 1-hop intermediate node list attached with text feature}}`, respectively. Will `({{candidate node}},{{candidate text feature}})` be connected to `({{central node}},{{text feature}})` within two hops through `{{the specified 1-hop intermediate node attached with text feature}}`?

Target template: `{{yes/no}}`

Prompt ID: 2-2-2-3

Input template:

{{central node}}, {{text feature}} is connected with {{2-hop neighbor list attached with text feature}} within two hops. Which other node will be connected to {{central node}}, {{text feature}} within two hops?

Target template: {{node_id}}

Prompt ID: 2-2-2-4

Input template:

{{central node}}, {{text feature}} is connected with {{2-hop neighbor list attached with text feature}} within two hops through {{the corresponding 1-hop intermediate node list attached with text feature}}, respectively. Which other node will be connected to {{central node}}, {{text feature}} within two hops through {{the specified 1-hop intermediate node attached with text feature}}?

Target template: {{node_id}}

Prompt ID: 2-2-3-1

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops. Will {{candidate node}}, {{candidate text feature}} be connected with {{central node}}, {{text feature}} within three hops?

Target template: {{yes/no}}

Prompt ID: 2-2-3-2

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops through {{the corresponding 2-hop intermediate path list attached with text feature}}, respectively. Will {{candidate node}}, {{candidate text feature}} be connected to {{central node}}, {{text feature}} within three hops through {{the specified 2-hop intermediate path attached with text feature}}?

Target template: {{yes/no}}

Prompt ID: 2-2-3-3

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops. Which other node will be connected to {{central node}}, {{text feature}} within three hops?

Target template: {{node_id}}

Prompt ID: 2-2-3-4

Input template:

{{central node}}, {{text feature}} is connected with {{3-hop neighbor list attached with text feature}} within three hops through {{the corresponding 2-hop intermediate path list attached with text feature}}, respectively. Which other node will be connected to {{central node}}, {{text feature}} within three hops through {{the specified 2-hop intermediate path attached with text feature}}?

Target template: {{node_id}}