

LLM-Empowered Class Imbalanced Graph Prompt Learning for Online Drug Trafficking Detection

Tianyi Ma¹, Yiyue Qian^{2†}, Zehong Wang¹, Zheyuan Zhang¹,
Chuxu Zhang³, Yanfang Ye^{1*}

¹ University of Notre Dame, Notre Dame, IN, USA

² Amazon GenAI, Seattle, WA, USA

³ University of Connecticut, Storrs, CT, USA

tma2@nd.edu, yyqian5@gmail.com, {zwang43, z Zhang42}@nd.edu,
chuxu.zhang@uconn.edu, yye7@nd.edu

Abstract

As the market for illicit drugs remains extremely profitable, major online platforms have become direct-to-consumer intermediaries for illicit drug trafficking participants. These online activities raise significant social concerns that require immediate action. Existing approaches to combating this challenge are generally impractical due to the imbalance of classes and scarcity of labeled samples in real-world applications. To this end, we propose a novel Large Language Model-empowered Heterogeneous Graph Prompt Learning framework for illicit Drug Trafficking detection, called LLM-HetGDT, that leverages LLM to facilitate heterogeneous graph neural networks (HGNNs) to effectively identify drug trafficking activities in the class-imbalanced scenarios. Specifically, we first pre-train HGNN over a contrastive pretext task to capture the inherent node and structure information over the unlabeled drug trafficking heterogeneous graph (HG). Afterward, we employ LLM to augment the HG by generating high-quality synthetic user nodes in minority classes. Then, we fine-tune the soft prompts on the augmented HG to capture the important information in the minority classes for the downstream drug trafficking detection task. To comprehensively study online illicit drug trafficking activities, we collect a new HG dataset over Twitter, called Twitter-HetDrug. Extensive experiments on this dataset demonstrate the effectiveness, efficiency, and applicability of LLM-HetGDT. Our source code and dataset are available [here](#).

1 Introduction

As the market for illicit drugs (e.g., methamphetamine and heroin) remains immensely profitable, the crime of drug trafficking has continually evolved with modern technologies, e.g., online platforms. Recent studies (Qian et al., 2021b; Ma et al.,

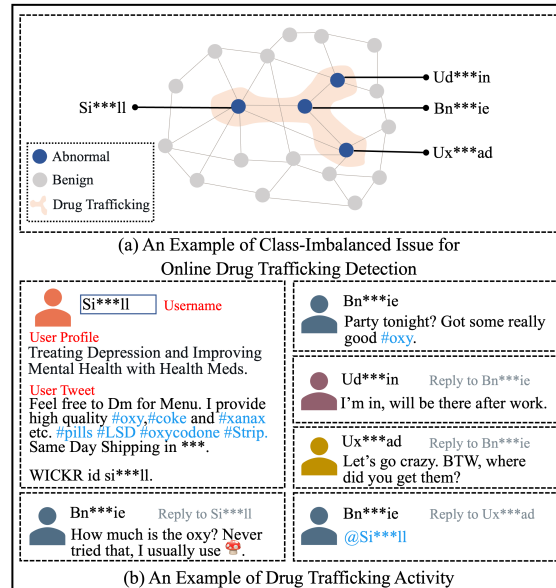


Figure 1: Illustration of drug trafficking activities among users on Twitter.

2023; Zhang et al., 2020; Wen et al., 2022) have demonstrated that the major online platforms, e.g., Twitter/X, Facebook, and Instagram, have become direct-to-consumer intermediaries for illicit drug trafficking. As depicted in Figure 1 (b), illicit drug trafficking participants can effortlessly create accounts on these platforms to promote their products through drug-related posts and quickly attract other abnormal users to share or trade drugs on social media platforms. Consequently, these drug-related behaviors among users naturally form drug trafficking networks on social media platforms, and these illegal drug-related activities have turned into a societal concern due to their catastrophic consequences, from violent crimes to public health (e.g., over 107,941 drug overdose deaths occurred in 2022 in the U.S. (Centers for Disease Control and Prevention, 2023)). To combat illicit drug trafficking online, there is an imminent need to develop novel and effective techniques for detecting drug trafficking activities on social media.

To address this challenge, existing studies can

*Corresponding author.

†The work is not related to the position at Amazon.

be categorized into two approaches: (i) large language models (LLMs)-based frameworks (Hu et al., 2024, 2023; Qian et al., 2025c) and (ii) graph-based methods (Ma et al., 2023; Qian, 2024; Wang et al., 2024; Ma et al., 2025b; Wang et al., 2025d). However, these two approaches still face limitations, such as class-imbalanced and label-scarcity issues. On one hand, LLM-based methods may struggle to efficiently identify drug trafficking activities on social networks because of the latency limitation. This inefficiency is exacerbated due to the class-imbalanced issue. Specifically, as illustrated in Figure 1 (a), merely four out of twenty users are involved in drug trafficking activities, and LLM-based methods require a significant amount of time to identify drug trafficking participants among huge social networks, making it impractical in real-world applications. Moreover, LLM-based methods primarily utilize user-level information, such as user profiles and posts for drug trafficking detection, while ignoring the rich and complex relationship information, e.g., users following others and replying to posts, etc. For example, Hu et al. (2024) utilizes anchor posts to identify drug-related users on Instagram, while ignoring contextual relationships between users. As shown in Figure 1 (b), it is challenging to classify the users "Ud***in" and "Ux***ad" by analyzing their posts. However, by examining the rich contextual relationships between users, such as their replies to the known drug buyer "Bn***id", we can successfully identify them as drug users. On the other hand, graph-based methods typically study drug trafficking in balanced data settings, which fail to consider inherent class-imbalanced issues, making it difficult for real-world scenarios. In addition, they require sufficient labeled samples for training, which is time and labor-intensive, and ignore the rich information in the unlabeled data. This naturally raises the following question: *How can we design an effective and efficient framework to detect drug-related users on social networks?*

To this end, we design a novel Large Language Model-enhanced Heterogeneous Graph Prompt Learning framework called **LLM-HetGDT** for Drug Trafficking detection. Specifically, we leverage heterogeneous graphs (HG) to model the complex relationships among entities, i.e., users, posts, and keywords, and further introduce the "pre-train, prompt" paradigm over heterogeneous graph neural networks (HGNNs) to comprehensively study drug trafficking activities on social media. In the pre-

train stage, we train an HGNN for the contrastive pretext task to capture the rich structure information in unlabeled data. Afterward, inspired by the success of LLMs in natural language understanding and synthetic data generation (Ma et al., 2025a; Tang et al., 2023; Ma et al., 2025c; Long et al., 2024), we leverage LLMs to augment the original HG by generating synthetic user nodes in minority classes, i.e., drug trafficking participants, enabling the model to learn valuable information from drug trafficking participant nodes. We further fine-tune three types of prompts, i.e., node, structure, and drug trafficking prompts on the augmented HG to facilitate the pre-trained HGNNs to learn the node representations for downstream tasks. To conclude, our work makes the following contributions:

- **Novelty:** We develop a novel framework called LLM-HetGDT that integrates HGNNs with LLMs for heterogeneous graph representation learning to effectively address class imbalance and label scarcity issues in illicit drug trafficking detection. To the best of our knowledge, this is the **first work** that employs heterogeneous graph prompt learning and LLMs to study online drug trafficking activities with limited labeled samples under class-imbalanced settings.
- **New Data:** To comprehensively study drug trafficking activities, we collect a new drug trafficking dataset from Twitter and further construct a heterogeneous graph called Twitter-HetDrug, which contributes to research communities of drug trafficking and (class-imbalanced) graph learning.
- **Effectiveness:** Comprehensive experiments on Twitter-HetDrug demonstrate the effectiveness of our proposed method in comparison with state-of-the-art baseline methods.

2 Backgrounds

Preliminary. A heterogeneous graph is defined as a network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{A}, \mathcal{R})$, where \mathcal{V} is a node set with node type set \mathcal{A} and \mathcal{E} is the edge set with relation type set \mathcal{R} . Here, $\mathcal{X} = \{\mathbf{X}^A | A \in \mathcal{A}\}$ represents a set of the node attribute feature matrix $\mathbf{X}^A \in \mathbb{R}^{N_A \times d_A}$ with node size N_A and dimension d_A for node type $A \in \mathcal{A}$. The target node set \mathcal{V}^T is a set of nodes from node type $T \in \mathcal{A}$ for node classification. The class imbalance ratio $\text{CIR}_{i,j} = |\mathcal{V}_{C_i}^T|/|\mathcal{V}_{C_j}^T|$, where $\mathcal{V}_{C_i}^T$ and $\mathcal{V}_{C_j}^T$ denote

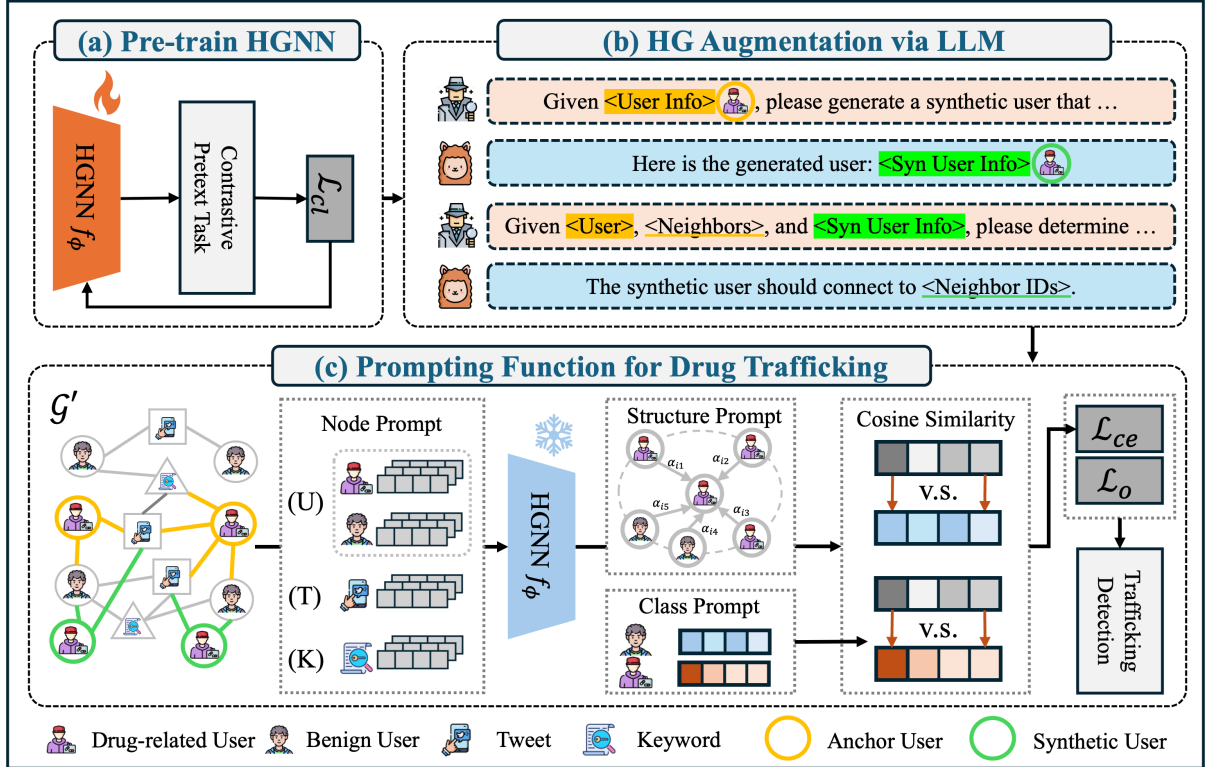


Figure 2: The overall framework of LLM-HetGDT: (a) It pre-trains HGNN with the contrastive pretext task. (b) It leverages LLM to generate synthetic users and connections between synthetic users and neighbors of the original users, forming an augmented HG \mathcal{G}' ; (c) LLM-HetGDT injects node prompt to node attribute features and feeds the augmented HG \mathcal{G}' into the pre-trained HGNN to obtain the target node embeddings. Afterward, it augments the target node embeddings with structure prompt and further computes the similarity between node embeddings and class prompt to obtain the classification loss for optimization.

the set of nodes belongs to class C_i and C_j , respectively. If $\text{CIR}_{i,j} > 1$, then C_i is called the majority class, and C_j is called the minority class, vice versa.

Related work. We provide relevant works to tackle the class imbalance issues in this section. Existing works against the class imbalanced issue in graphs can be roughly divided into three categories, i.e., post-hoc correction (Kang et al., 2020; Tian et al., 2020), loss modification (Mohammed et al., 2020; Kang et al., 2019), and re-sampling (Chawla et al., 2002; Qian et al., 2025a). For example, GraphSmote (Zhao et al., 2021) and ImGAGN (Qu et al., 2021) introduce SMOTE (Chawla et al., 2002) and GAN (Goodfellow et al., 2020) in graph augmentation, respectively, to generate synthetic minority nodes while learning the topological structure distribution at the same time (Gao et al., 2023). DR-GCN (Shi et al., 2020) and RA-GCN (Ghorbani et al., 2022) propose re-weighting mechanisms to balance the loss of minority and majority classes. Motivated by existing re-sampling techniques, we propose an oversampling framework in the prompt tuning stage to obtain more effective prompts for

drug trafficking detection tasks. The previous methods primarily over-sample synthetic nodes from the attribute feature level, potentially introducing redundant or harmful information. In light of this, this work employs large language models to generate synthetic nodes from the text attribute level to obtain more accurate and valuable synthetic nodes. Additional related work for graph-based methods is provided in Appendix A.

Problem Definition. *Class-imbalanced Graph Prompt Learning for Drug Trafficking Detection.* Given a class-imbalanced heterogenous graph $\mathcal{G} = (\mathcal{V}, \mathcal{X}, \mathcal{E}, \mathcal{A}, \mathcal{R})$ built on real-world drug trafficking data over social media, the objective is to build a graph prompt learning model to detect drug trafficking participants.

3 Methodology

In this section, we present the details of LLM-HetGDT, which includes two key steps: (i) pre-train HGNN, and (ii) LLMs-enhanced prompt tuning for drug trafficking detection. Mention that the original HG contains text-attributed features, and we leverage SentenceBert (Reimers and Gurevych,

2019) to encode text attributes $\mathcal{T} = \{t_i | v_i \in \mathcal{V}\}$ into attribute features $\mathcal{X} = \{x_i | v_i \in \mathcal{V}\}$. The overall framework is illustrated in Fig. 2.

3.1 Pre-train HGNN

As we mentioned above, previous studies face the label-scarcity issue, and underestimate the rich information in the unlabeled data. Therefore, inspired by the success of "pre-train, prompt" in natural language processing (Liu et al., 2023a), we propose a novel LLM-enhanced HG prompt learning framework that pre-trains HGNN via contrastive pretext tasks to capture rich contextual and structure information in HG, and further fine-tune prompts to reformulate the downstream detection task better align with the pretext task.

Pre-training Stage. Our framework is applicable to any HGNN. In this work, we utilize a cross-view contrastive learning method (Wang et al., 2021, 2025b), as an example to map target nodes \mathcal{V}^T into low-dimensional representations $\mathbf{Z} \in \mathbb{R}^{N_T \times d}$, where N_T is the size of target nodes and d is the latent dimension. Given an HG $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{A}, \mathcal{R})$ that models drug trafficking activities on social media, we first adopt a type-specific mapping layer to project all the node attribute features \mathcal{X} into the same space: $\tilde{\mathbf{X}}^A = \mathbf{X}^A W^A + b^A$, where $\mathbf{X}^A \in \mathbb{R}^{N_A \times d_A}$, $W^A \in \mathbb{R}^{d_A \times d}$, and $b^A \in \mathbb{R}^d$ denote the attribute feature matrix, projection matrix, and the bias vector for node type $A \in \mathcal{A}$, respectively. Afterward, we employ meta-path view and network schema view encoders (Wang et al., 2021) to map each target node $v_i \in \mathcal{V}^T$ into two embeddings $z_i^{mp} \in \mathbb{R}^d$ and $z_i^{sc} \in \mathbb{R}^d$, respectively. These two embeddings are then fed into the pretext objective function to optimize the HGNN.

Pre-train Optimization. Inspired by GCL (You et al., 2020; Liu et al., 2022; Wang et al., 2025f; Wen et al., 2024; Qian et al., 2022d, 2024, 2025b), we leverage the temperature-scaled contrastive loss as the objective function for the pretext task. The objective function is formulated as follows:

$$\mathcal{L}_{cl} = -\frac{1}{N} \sum_{v_i \in \mathcal{V}^T} \frac{\exp(\delta_{i,i}/\tau)}{\sum_j \exp(\delta_{i,j}/\tau)}, \quad (1)$$

where $\delta_{i,j}$ is the cosine similarity between z_i^{mp} and z_j^{sc} , τ is the temperature hyperparameter. We adopt z^{mp} as the target node representations \mathbf{Z} since the target type nodes explicitly participate in aggregating z^{mp} . With the pre-trained HGNN, we further devise an LLM-enhanced prompt tuning framework for drug trafficking detection tasks.

3.2 Prompt Tuning for Drug Trafficking Detection

3.2.1 HG Augmentation via LLM

As drug trafficking participants are inactive and rare on social media, the HG to model real-world online drug trafficking activities is a class-imbalanced network. This class imbalance issue may lead to a performance bias, as the majority class tends to outperform the minority class due to their more significant number of labeled nodes (Zhao et al., 2021; Shi et al., 2020; Liu et al., 2023b). To tackle this issue, existing works (Chawla et al., 2002; Park et al., 2022; Zhao et al., 2021) mainly balance the data distribution through over-sampling techniques from the attribute feature level. However, due to the limited number of nodes in the minority class and the inevitable information loss in feature extraction from raw textual information to attribute features (Li et al., 2020), the sampled synthetic nodes may introduce redundant or even harmful information for downstream tasks.

Synthetic User Generation. To handle the above challenge, inspired by the success of LLMs in generating synthetic data (Tang et al., 2023; Long et al., 2024), we leverage an LLM to generate synthetic minority users in natural languages, thereby balancing the data distribution toward the minority class (drug trafficking participants). Specifically, for each labeled user node v_i in the minority class, we feed the user text content t_i into LLM to obtain a synthetic node v'_i with corresponding user text content t'_i . Formally, we have:

$$t'_i = \text{LLM}(\text{Prompt}_{\text{node}}(t_i)), \quad (2)$$

where user text content t_i contains the user i 's information, including username, user ID, and user profile. Afterward, we employ SentenceBert over user text content t'_i to obtain the attribute feature x'_i for synthetic node v'_i :

$$x'_i = \text{SentenceBert}(t'_i). \quad (3)$$

The prompt template for synthetic user generation is listed in Appendix Table 8.

Edge Generation Among Synthetic Nodes. Up to now, we have obtained a set of synthetic nodes corresponding to the labeled user nodes in the minority class. However, these nodes are still isolated in the HG \mathcal{G} . Intuitively, for each pair of labeled user node v_i and its corresponding synthetic node v'_i , the synthetic node v'_i should have connections

with the neighbors of the labeled node v_i . However, connecting the synthetic node v'_i with every neighbor of node v_i is not feasible, as the synthetic node contains different information than the original node. Therefore, to select more appropriate and informative neighbors for synthetic nodes, we further employ LLM to generate edges between synthetic node v_i and the neighbors of the original node v_i , i.e., $\mathcal{N}(v_i)$. Mathematically, we have:

$$e = \text{LLM}(\text{Prompt}_{\text{edge}}(t'_i, t_j)). \quad (4)$$

The prompt template for edge generation is listed in Appendix Table 9. Through aforementioned steps, we have an augmented HG $\mathcal{G}' = \{\mathcal{V}', \mathcal{E}', \mathcal{X}', \mathcal{A}, \mathcal{R}\}$, where $\mathcal{V}' = \mathcal{V} \cup \{v'_i | v_i \in \mathcal{V}_c\}$, $\mathcal{E}' = \mathcal{E} \cup \{(v'_i, v_j) | v_j \in \mathcal{N}^A(v'_i), A \in \mathcal{A}\}$, and $\mathcal{X}' = \mathcal{X} \cup \{x'_i | v_i \in \mathcal{V}_C\}$. Here, \mathcal{V}_C is a set of nodes in minority class $C \in \mathcal{C}$, $\mathcal{N}^A(v_i)$ denotes the neighbors of node v_i in node type A . Notice that we merely sample synthetic nodes for labeled user nodes in the minority class. Moreover, in the pre-train stage, as the label information is unavailable, we only employ the original HG \mathcal{G} for pre-training. The augmented HG \mathcal{G}' is exclusively used to tune prompts to learn valuable information among drug trafficking participants.

3.2.2 HG Prompting Function Design

As the training objective gap between the pretext contrastive task and the downstream task, the pre-trained HGNN may fail to identify the drug trafficking participants effectively. Hence, we introduce a novel prompting function, integrated with node, structure, and drug trafficking prompts, to bridge the training objective gap between the pretext task and the downstream drug trafficking detection task. **Node Prompt.** Previous studies (Fang et al., 2024; Ma et al., 2024; Sun et al., 2023) have demonstrated that incorporating a small number of trainable parameters directly into the node feature space can effectively tailor pre-trained node representations, facilitating the knowledge transfer from the pretext task to the downstream task. However, existing methods primarily introduce feature weight-based prompt (Liu et al., 2023c; Yu et al., 2024) or learnable independent basis tokens (Sun et al., 2022; Ma et al., 2024) as prompts, which the majority class samples can simply dominate the prompt tuning process. To tackle this issue, besides the LLM-enhanced over-sampling technique in Section 3.2.1, we devise a set of independent basic

tokens for each class, serving as class-specific augmentations to the original target node attribute features. Specifically, for each class $C \in \mathcal{C}$, the class-specific feature tokens $\mathbf{F}^C \in \mathbb{R}^{K \times d_T}$ are defined as: $\mathbf{F}^C = [\mathbf{f}_1^C, \mathbf{f}_2^C, \dots, \mathbf{f}_K^C]^\top$, where $\mathbf{f}_k^C \in \mathbb{R}^{d_T}$, d_T is the attribute feature dimension for target node type T and K is a hyper-parameter. Afterward, the feature tokens \mathbf{F}^C are injected into the target node attribute features \mathbf{X} :

$$\bar{\mathbf{F}}^C = \text{Att}(\mathbf{X}, \mathbf{F}^C), \quad (5)$$

$$\hat{\mathbf{X}}_{i,:} = \mathbf{X}_{i,:} + \text{PMA}(\{\bar{\mathbf{F}}^C | C \in \mathcal{C}\}), \quad (6)$$

where $\bar{\mathbf{F}}^C \in \mathbb{R}^{d_T}$ is the linear combination of feature tokens \mathbf{F}^C through an attention mechanism $\text{Att}(\cdot)$ and $\text{PMA}(\cdot)$ is the set transformer function (Lee et al., 2019). As attention and set transformer mechanisms are not the main contribution of this paper, we refer readers to Appendix D.1 and D.2 for more details about these two mechanisms. The key insight behind the class-specific feature tokens is to separate the prompt for each class to alleviate the influence of class imbalance and further capture the unique characteristics of each class equally. Besides, considering the heterogeneity of HGs, we also employ a set of tokens for the rest of the node types $A \in \mathcal{A}/T$: $\mathbf{F}^A = [\mathbf{f}_1^A, \mathbf{f}_2^A, \dots, \mathbf{f}_K^A]^\top$, where $\mathbf{f}_k^A \in \mathbb{R}^{d_A}$, d_A is the attribute feature dimension for node type A . Then, we reuse the attention mechanism $\text{Att}(\cdot)$, and inject the prompt into the attribute feature \mathbf{X}^A :

$$\hat{\mathbf{X}}_{i,:}^A = \mathbf{X}_{i,:}^A + \text{Att}(\mathbf{X}_{i,:}^A, \mathbf{F}^A). \quad (7)$$

Subsequently, the prompted node features $\hat{\mathcal{X}} = \{\hat{\mathbf{X}}^A | A \in \mathcal{A}\}$, and the augmented HG \mathcal{G}' are fed into pre-trained HGNN to obtain the target node embeddings $\mathbf{Z} \in \mathbb{R}^{N_T \times d}$.

Structure Prompt. Merely relying on the feature prompt may not be sufficient to transfer the learned relation information from pre-training to drug trafficking detection tasks. Prior studies (Ma et al., 2023; Qian et al., 2021b) have demonstrated that drug trafficking participants are more likely to connect with other participants on social media. Therefore, we propose a structure prompt to leverage the expressive meta-path based neighborhood information for robust classification performance. Particularly, for each target node v_i , we employ the set transformer function $\text{PMA}(\cdot)$ to aggregate the neighbor node embeddings with respect to meta-path $P \in \mathcal{P}$:

$$\mathbf{H}_{i,:}^P = \text{PMA}(\{\mathbf{Z}_{j,:} | v_j \in \mathcal{N}^P(v_i)\}). \quad (8)$$

Here a meta-path P is a path defined in the form of $A_1 \xrightarrow{R_1} A_2 \xrightarrow{R_2} \dots \xrightarrow{R_l} A_{l+1}$, where $A_i \in \mathcal{A}$ and $R_i \in \mathcal{R}$. It describes a composite relation $R = R_1 \circ R_2 \circ \dots \circ R_l$ between nodes A_1 and A_{l+1} , where \circ denotes the composition operator. Afterward, we employ semantic attention (Wang et al., 2019), i.e., **SM-Att**(\cdot), to aggregate the structure embeddings as structure tokens \mathbf{S} among meta-paths \mathcal{P} :

$$\mathbf{S}_{i,:} = \mathbf{SM-Att}(\{\mathbf{H}_{i,:}^P | P \in \mathcal{P}\}). \quad (9)$$

Again, we leave the detailed formulation of meta-path \mathcal{P} and semantic attention to Appendix D.3 and D.4, respectively.

Drug Trafficking Prompt. Instead of the discrete class labels for drug trafficking participants and benign, we introduce the drug trafficking prompt, which contains a set of class prototype tokens, i.e., $\mathbf{C} = [\mathbf{p}, \mathbf{n}]$, where $\mathbf{C} \in \mathbb{R}^{2 \times d}$, and d is the same dimension as the latent node embeddings. Here, \mathbf{p} and \mathbf{n} denote the prototype tokens for drug trafficking participants and benign, respectively. Moreover, each prototype token is initialized by the mean of labeled nodes embeddings:

$$\mathbf{p} = \frac{1}{|\mathcal{V}_p|} \sum_{v_i \in \mathcal{V}_p} \mathbf{z}_{i,:}, \mathbf{n} = \frac{1}{|\mathcal{V}_n|} \sum_{v_i \in \mathcal{V}_n} \mathbf{z}_{i,:}. \quad (10)$$

These task tokens serve as class prototypes for node classification, conveying the discrete labels to class-specific semantics (Sun et al., 2023, 2022; Ma et al., 2024). The key idea behind the task prompt is shifting the node classification to pairwise similarity calculation between the node embeddings and prototype tokens, which can seamlessly align with the pretext contrastive task. Moreover, the prompt tuning process facilitates the prototype tokens to represent the unique characteristics of each class, even some have fewer labeled samples.

3.2.3 Prompt Tuning

With the drug trafficking prompts \mathbf{C} , structure prompt \mathbf{S} , and target node embeddings \mathbf{Z} , we employ an MLP to project them into the sample embedding space. Moreover, the structure prompt is injected into the target node embeddings to obtain the final node embeddings: $\mathbf{Z}' = \mathbf{Z} + \delta \mathbf{S}$, where δ is a hyperparameter. To align with the pre-training stage, we adopt the same temperature-scaled contrastive loss in Equation 1 as the objective function

for node classification loss:

$$\mathcal{L}_{ce} = -\frac{1}{N} \sum_{v_i \in \mathcal{V}_T} \frac{\exp(\xi_{i,y_i}/\tau)}{\exp(\xi_{i,p}/\tau) + \exp(\xi_{i,n}/\tau)}, \quad (11)$$

where ξ_{i,y_i} is the cosine similarity between the prompted node embeddings $\mathbf{Z}'_{i,:}$ and the class prompt corresponding to ground-truth label $y_i \in \{\mathbf{p}, \mathbf{n}\}$. Moreover, we employ orthogonal constraints to keep the drug trafficking prompt orthogonal during prompt tuning:

$$\mathcal{L}_o = \|\mathbf{C}\mathbf{C}^\top - \mathbf{I}\|_F^2, \quad (12)$$

where \mathbf{I} is the identity matrix. Then, the final objective for prompt tuning is formulated as:

$$\mathcal{L}_{pt} = \mathcal{L}_{ce} + \lambda \mathcal{L}_o, \quad (13)$$

where λ is the trade-off hyperparameter. The pseudocode for the training procedure is provided in Appendix Alg. 1.

4 Experiments

4.1 New Dataset for Drug Trafficking Detect

To comprehensively study online drug trafficking activities in a real-world scenario, we build a new real-world HG dataset on Twitter, called **Twitter-HetDrug**, to analyze drug trafficking activities under a class-imbalanced setting, i.e., merely 11.9% of users are drug trafficking participants. Our data is collected through the official Twitter API from Dec. 2020 to Aug. 2021. More detailed discussions about the dataset (i.e., data collection, construction and annotation) and data statistics are provided in Appendix B and Appendix Table 6, respectively.

4.2 Experimental Setup

Baseline Methods. To evaluate LLM-HetGDT, we compare it with Twenty-one state-of-the-art (SOTA) baseline methods, which are divided into seven groups: graph representation learning methods (G1), heterogeneous graph representation learning models (G2), heterogeneous graph self-supervised learning methods (G3), generic methods against class imbalance data (G4), graph models against the class imbalance issue (G5), graph prompt learning models (G6), and LLM-based methods (G7). A detailed discussion about baseline methods is provided in Appendix E. Due to the limited space, we only list the best model performance of each group in Table 1, and the detailed performance with standard deviation of each baseline method can be found in Appendix H.

Setting		Twitter-HetDrug 10%		Twitter-HetDrug 20%		Twitter-HetDrug 40%	
Group	Model	Macro-F1	GMean	Macro-F1	GMean	Macro-F1	GMean
G1	GCN	55.74	53.92	58.38	56.43	60.94	59.72
G2	HGT	60.95	59.88	63.85	63.70	67.08	66.37
G3	HeCo	62.93	60.85	64.83	63.21	67.94	66.52
G4	Over-sampling	63.79	62.27	65.10	63.08	68.35	67.39
G5	PC-GNN	65.98	63.84	67.52	64.71	70.59	71.94
G6	HGPrompt	66.85	64.37	67.97	66.08	69.72	69.42
G7	KG-GPT	65.74	63.92	65.36	63.98	65.18	63.65
Ours	LLM-HetGDT	70.52	71.69	71.36	73.87	74.06	76.24

Table 1: Performance comparison among various methods on Twitter-HetDrug for drug trafficking detection. Bolded numbers indicate the best performance, and underlined numbers represent runner-up performance.

Experimental Settings. All experiments are conducted in the environment of the Ubuntu 22.04 OS, plus Intel i9-12900K CPU, four Nvidia A40 Graphic Cards, and 48 GB of RAM. We train all methods with a fixed number of epochs (i.e., 500). We use Llama3.1 70B (team, 2024) as our LLM backbone. Besides, all supervised methods, i.e., in G1-G6, are trained five times, and the average performance on testing data is reported. We utilize Adam (Kingma and Ba, 2015) as the optimizer and run a grid search on hyperparameters for each model. Then, we report the best performance among the optimal hyperparameters for each method. Detailed discussion for experimental settings is provided in Appendix F. For LLM-based methods, we set the temperature to 0 to ensure the deterministic output. Following existing works (Liu et al., 2021; Park et al., 2022; Goodfellow et al., 2020) in evaluating class imbalance classification, we adopt two metrics, i.e., Macro-F1 score (**Macro-F1**) and GMean score (**GMean**), to evaluate all models.

4.3 Performance Comparison

Table 1 lists the performance of selected baseline methods and LLM-HetGDT over Twitter-HetDrug with different percentages, i.e., 10%, 20%, and 40%, of training data for drug trafficking detection. Besides, we use 10% samples for validation and the rest for testing. According to Table 1, we can conclude that: (i) HG methods outperform graph methods, i.e., all heterogeneous graph representation learning models in G2 and G3 show better performance than graph-based methods in G1. This result indicates **the necessity of heterogeneous graphs for drug trafficking detection tasks**. (ii) Class-imbalanced methods in G4 and G5 outperform all (heterogeneous) graph representation learning methods in G1-G3, demonstrating

that **class imbalance is a significant challenge for online drug trafficking detection tasks**, where traditional semi-supervised (heterogeneous) graph representation learning may struggle to learn the patterns from minority classes. (iii) HGPrompt is the best baseline model in low-labeled settings, i.e., 10% and 20%, however in high-labeled settings, the best baseline model is PC-GNN. This indicates the superiority of graph prompt learning in low-labeled settings, while methods that are specifically designed for class-imbalanced issues in G4 and G5 may fail to model the minority classes due to the limited labeled data. However, in high-labeled settings, methods specifically designed for class-imbalanced data may learn enough information from minority classes. This validates **the need to design effective techniques for class imbalance issues in low-label settings**. (v) LLM-based methods, i.e., in G7, have consistent performance in different label settings, as we do not fine-tune LLMs, and merely prompt it for prediction. Despite this, LLM-based approaches still outperform some supervised methods in G1-G4, demonstrating the superiority of natural language understanding and generation ability of LLMs. This is one of the reasons we propose to leverage LLMs to generate synthetic user nodes to tackle class-imbalance issue in drug trafficking detection tasks. (vi) Our model LLM-HetGDT gains the best performance by comparison with all baseline models, which demonstrates **the effectiveness of our model over Twitter-HetDrug**.

4.4 Ablation Study

4.4.1 Prompt Component Analysis

To show the effectiveness of each prompt in LLM-HetGDT, we further analyze the contribution of each component, i.e., entire prompts (A1), drug trafficking prompt (A2), node prompt (A3), and




Figure 3: Performances of model variants over Twitter-HetDrug with training splits 10%, 20%, and 40%.

Class	# nodes	PC-GNN	HGPrompt	Ours
Neg.	25,402	88.39	87.08	90.47
Pos.	3,437	43.57	46.62	50.65
Overall	28,839	65.98	66.85	70.52

Table 2: Macro-F1 among models over Twitter-HetDrug-10%. Pos. and Neg. indicates drug participants and benign users, respectively.

structure prompt (A4), by removing it separately, as illustrated in Fig 3. First, we remove all three prompts, and the significant drop in performance indicates the effectiveness of our proposed framework. Afterward, we remove the drug trafficking prompt (A2), node prompt (A3), and structure prompt (A4), respectively. The performance of A2 decreases obviously, showing the necessity of the prototypes to capture the complex characteristics among drug trafficking participants for classification tasks. Moreover, the decline in the performance of A3 and A4 indicates the effectiveness of the node and structure prompts in our model.

4.4.2 Minority Class Analysis

To clearly show the improvement in the minority class, Table 2 lists Macro-F1 among some models for Twitter-HetDrug-10%. We find out that all models gain excellent performance over the majority class while they fail to achieve good performance over the minority class (participants). Our proposed method outperforms all baseline methods over both minority and majority classes, which again demonstrates the effectiveness of our model.

4.5 Deployment Evaluation

To validate the applicability of our model, we deploy LLM-HetGDT in a live manner to detect drug trafficking activities over Twitter in Jan. 2025, and compare with two SOTA drug trafficking detection models, i.e., MetaHG (Qian et al., 2021b) and KG-GPT (Hu et al., 2024). The detailed settings of the deployment are provided in Appendix G, and the results are shown in Table 3. According to the table, we conclude that our proposed method outperforms the baseline methods in every aspect,

which indicates the effectiveness, efficiency, and applicability of our model for online illicit drug trafficking detection tasks.

Model	Macro-F1 ↑	GMean ↑	Avg. Time ↓
MetaHG	61.73	62.48	0.39s
KG-GPT	64.24	63.18	4.72s
Ours	66.38	65.35	0.18s

Table 3: The performance of drug trafficking detection methods for deployment evaluation.

4.6 Case Study

We analyze the synthetic users generated by LLM. Table 4 illustrates an example of an input prompt with original user context and the output synthetic user generated by LLM. According to Table 4, the original user profile contains three key phrases, i.e., "online dispensary," "premium lab-tested psychedelic products," and hashtags "#LSD #420 #acid" (highlighted in green, red, and blue, respectively), that help to identify the user "yas***exb" as a drug trafficking participant. In the output, the LLM effectively captures these key phrases that best depict the synthetic user and introduce new information, i.e., "Get ready to ..." and street name hashtags to ensure variability, which aligns with our prompt design. We conduct another case study for edge generation in Appendix J.

5 Conclusion

In this study, to handle the class-imbalanced and label-scarcity issues in existing illicit drug trafficking detection methods, we propose a novel LLM-enhanced heterogeneous graph prompt learning framework, called LLM-HetGDT, to detect online drug trafficking participants in real-world scenarios. Specifically, we leverage LLMs to augment the heterogeneous graph (HG) by generating synthetic nodes in minority class, i.e., drug trafficking participants, facilitating the heterogeneous graph neural networks (HGNNs) to learn valuable information in minority class. To comprehensively study the drug trafficking activities on social media, we collect a new HG dataset over Twitter, called Twitter-HetDrug. Extensive experiments on the dataset demonstrate the effectiveness, efficiency, and applicability of our framework.

6 Limitations

In our work, we follow existing studies (Ma et al., 2023; Qian et al., 2021b) that utilize SentenceBert

INPUT:

Instruction: <The same instruction in Table 8>.

Context:

Username: yas***exb;

User ID: @dru***cy;

User Profile: **Online dispensary, premium lab-tested psychedelic products**, must be 18+, same day shipping in ***. DM me for menu. **#LSD #420 #Meth #acid**.

OUTPUT:

Username: Psy***ete;

User ID: @ss***ete;

User Profile: **Your one-stop shop for lab-tested psychedelic products**. Get ready to trip with the selection of **premium psychedelic products**. **#lsdtrip #psychedelics #acid #meth #420**.

Table 4: Example of synthetic user generated by LLM.

to encode the text attributes. We acknowledge that there are alternatives available in the literature that may have better performance. Moreover, we leverage open-source LLMs, i.e., Llama3.1 70B, to generate synthetic users and edges, which may have a lower capability compared to other LLMs, e.g., GPT-4o, DeepSeek-R1, etc. Despite this, our experimental results are still highly desirable. It remains uncertain whether more powerful text encoders or LLMs can further improve the performance of our proposed method to some extent. We regard these two limitations as promising directions for future work.

7 Ethical Considerations

Our method serves as a powerful tool for online drug trafficking detection. However, there is potential for the generated content to include biased or harmful information. To mitigate this risk, more careful prompt design, including clear instructions and guidelines, can help steer the LLMs toward generating positive and accurate content. Additionally, users must be mindful of ethical concerns such as bias in the data, and ensure responsible use of the tool in different applications.

8 Additional Discussion regarding Twitter-HetDrug

As a part of this study, we release our collected dataset, strictly following the ACL code of ethics and X (formerly Twitter) developer agreement. All confidential information, such as usernames, locations, etc, has been **anonymized** to protect user privacy. Besides, we acknowledge that annotating user behavior may introduce false positives or

misclassification. We make every effort to ensure annotations are accurate and unbiased, as discussed in Appendix B.3. Nevertheless, the users should **acknowledge the inherent risk of annotation errors**, and we emphasize that researchers should **consider these limitations** when using the dataset. We would like to emphasize that Twitter-HetDrug is **strictly intended for research purposes** and must not be used for commercial or profit-driven activities. The dataset will be made available in two distinct versions: a feature-based version and a text-attributed version. The feature-based version contains only the nodes (including any generated synthetic nodes) represented by attribute features in **embedding spaces**, along with the graph structures. Specifically, each node in this version is associated with a single attribute vector $x \in \mathbb{R}^{d_A}$. The purpose of the feature-based version is to facilitate the reproducibility of our experiments. **Since the attribute features are provided as embeddings, user confidentiality is preserved.** Access to the text-attributed version, which includes the text content, requires individuals to contact the authors. Recipients must **agree to the data usage agreement** and obtain **written permission from X (formerly Twitter)** before receiving the dataset.

9 Acknowledgements

This work was partially supported by the NSF under grants IIS-2321504, IIS-2334193, IIS-2340346, IIS-2217239, CNS-2426514, and CMMI-2146076, and Notre Dame Strategic Framework Research Grant (2025). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

- Rui Bing, Guan Yuan, Mu Zhu, Fanrong Meng, Huifang Ma, and Shaojie Qiao. 2023. Heterogeneous graph neural networks analysis: a survey of techniques, evaluations and applications. *Artificial Intelligence Review*.
- Centers for Disease Control and Prevention. 2023. Drug overdose deaths in the united states, 2002–2022.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. Smote: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In *CVPR*.

- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). Preprint, arXiv:2501.12948.
- Taoran Fang, Yunchao Zhang, Yang Yang, Chunping Wang, and Lei Chen. 2024. Universal prompt tuning for graph neural networks. In *NeurIPS*.
- Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. 2020. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *WWW*.
- Xinyi Gao, Wentao Zhang, Tong Chen, Junliang Yu, Hung Quoc Viet Nguyen, and Hongzhi Yin. 2023. Semantic-aware node synthesis for imbalanced heterogeneous information networks. In *CIKM*.
- Mahsa Ghorbani, Anees Kazi, Mahdieh Soleymani Baghshah, Hamid R Rabiee, and Nassir Navab. 2022. Ra-gcn: Graph convolutional network for disease prediction problems with imbalanced data. *Medical image analysis*.
- Chenghua Gong, Xiang Li, Jianxiang Yu, Cheng Yao, Jiaqi Tan, and Chengcheng Yu. 2024. Self-pro: A self-prompt and tuning framework for graph neural networks. *arXiv preprint arXiv:2310.10362*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Communications of the ACM*.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.
- Chuanbo Hu, Bin Liu, Xin Li, and Yanfang Ye. 2023. Unveiling the potential of knowledge-prompted chatpt for enhancing drug trafficking detection on social media. *arXiv preprint arXiv:2307.03699*.
- Chuanbo Hu, Bin Liu, Xin Li, Yanfang Ye, and Minglei Yin. 2024. Knowledge-prompted chatgpt: Enhancing drug trafficking detection on social media. *Information & Management*.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous graph transformer. In *WWW*.
- Hugging Face, Mistral AI. 2024. [Ministral-8b-instruct-2410](https://huggingface.co/mistralai/Ministral-8B-Instruct-2410). <https://huggingface.co/mistralai/Ministral-8B-Instruct-2410>.
- Xunqiang Jiang, Tianrui Jia, Yuan Fang, Chuan Shi, Zhe Lin, and Hui Wang. 2021. Pre-training on large-scale heterogeneous graph. In *KDD*.
- Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. 2019. Few-shot object detection via feature reweighting. In *ICCV*.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2020. Decoupling representation and classifier for long-tailed recognition. In *ICLR*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiosek, Seungjin Choi, and Yee Whye Teh. 2019. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. On the sentence embeddings from pre-trained language models. *arXiv preprint arXiv:2011.05864*.
- Nian Liu, Xiao Wang, Deyu Bo, Chuan Shi, and Jian Pei. 2022. Revisiting graph contrastive learning from the perspective of graph spectrum. In *NeurIPS*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023a. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*.
- Yang Liu, Xiang Ao, Zidi Qin, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. 2021. Pick and choose: a gnn-based imbalanced learning approach for fraud detection. In *WWW*.
- Zemin Liu, Yuan Li, Nan Chen, Qian Wang, Bryan Hooi, and Bingsheng He. 2023b. A survey of imbalanced learning on graphs: Problems, techniques, and future directions. *arXiv preprint arXiv:2308.13821*.
- Zemin Liu, Xingtong Yu, Yuan Fang, and Xinming Zhang. 2023c. Graphprompt: Unifying pre-training and downstream tasks for graph neural networks. In *WWW*.
- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. On llms-driven synthetic data generation, curation, and evaluation: A survey. *arXiv preprint arXiv:2406.15126*.
- Shang Ma, Tianyi Ma, Jiahao Liu, Wei Song, Zhenkai Liang, Xusheng Xiao, and Yanfang Ye. 2025a. Psyscam: A benchmark for psychological techniques in real-world scams. *arXiv preprint arXiv:2505.15017*.
- Tianyi Ma, Yiyue Qian, Chuxu Zhang, and Yanfang Ye. 2023. Hypergraph contrastive learning for drug trafficking community detection. In *ICDM*.
- Tianyi Ma, Yiyue Qian, Shinan Zhang, Chuxu Zhang, and Yanfang Ye. 2025b. Adaptive expansion for hypergraph learning. *arXiv preprint arXiv:2502.15564*.

- Tianyi Ma, Yiyue Qian, Zheyuan Zhang, Zehong Wang, Xiaoye Qian, Feifan Bai, Yifan Ding, Xuwei Luo, Shinan Zhang, Keerthiram Murugesan, et al. 2025c. Autodata: A multi-agent system for open web data collection. *arXiv preprint arXiv:2505.15859*.
- Yihong Ma, Ning Yan, Jiayu Li, Masood Mortazavi, and Nitesh V Chawla. 2024. Hetgpt: Harnessing the power of prompt tuning in pre-trained heterogeneous graph neural networks. In *WWW*.
- Roweida Mohammed, Jumanah Rawashdeh, and Malak Abdullah. 2020. Machine learning with oversampling and undersampling techniques: overview study and experimental results. In *ICICS*.
- Chanyoung Park, Donghyun Kim, Jiawei Han, and Hwanjo Yu. 2020. Unsupervised attributed multiplex network embedding. In *AAAI*.
- Joonhyung Park, Jaeyun Song, and Eunho Yang. 2022. GraphENS: Neighbor-aware ego network synthesis for class-imbalanced node classification. In *ICLR*.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*.
- Yiyue Qian. 2024. *Graph Representation Learning Techniques for the Combat Against Online Abusive Activity*. University of Notre Dame.
- Yiyue Qian, Philip Chen, Song Cui, and De Chen. 2023. Universal ring-of-abusers detection via multi-modal heterogeneous graph learning.
- Yiyue Qian, Tianyi Ma, Chuxu Zhang, and Yanfang Ye. 2024. Dual-level hypergraph contrastive learning with adaptive temperature enhancement. In *Companion Proceedings of the ACM Web Conference 2024*, pages 859–862.
- Yiyue Qian, Tianyi Ma, Chuxu Zhang, and Yanfang Ye. 2025a. Adaptive graph enhancement for imbalanced multi-relation graph learning. In *WSDM*.
- Yiyue Qian, Chuihui Zhang, Yiming Zhang, Qianlong Wen, Yanfang Ye, and Chuxu Zhang. 2022a. Modality imbalanced graph contrastive learning. In *NeurIPS*.
- Yiyue Qian, Chunhui Zhang, Yiming Zhang, Qianlong Wen, Yanfang Ye, and Chuxu Zhang. 2022b. Modality graph contrastive learning for imbalanced node classification. In *NeurIPS*.
- Yiyue Qian, Shinan Zhang, Lanhao Chen, Diego Socolinsky, Negin Sokhandan, Song Cui, De Chen, and Suchitra Sathyanarayana. 2025b. Enhancing e-commerce representation learning via hypergraph contrastive learning and interpretable llm-driven analysis. In *WWW*.
- Yiyue Qian, Shinan Zhang, Yun Zhou, Haibo Ding, Diego Socolinsky, and Yi Zhang. 2025c. Enhancing llm-as-a-judge via multi-agent collaboration.
- Yiyue Qian, Yiming Zhang, Nitesh Chawla, Yanfang Ye, and Chuxu Zhang. 2022c. Malicious repositories detection with adversarial heterogeneous graph contrastive learning. In *CIKM*.
- Yiyue Qian, Yiming Zhang, Qianlong Wen, Yanfang Ye, and Chuxu Zhang. 2022d. Rep2vec: Repository embedding via heterogeneous graph adversarial contrastive learning. In *SIGKDD*.
- Yiyue Qian, Yiming Zhang, Yanfang Ye, and Chuxu Zhang. 2021a. Adapting meta knowledge with heterogeneous information network for covid-19 themed malicious repository detection. In *IJCAI*.
- Yiyue Qian, Yiming Zhang, Yanfang Ye, and Chuxu Zhang. 2021b. Distilling meta knowledge on heterogeneous graph for illicit drug trafficker detection on social media. In *NeurIPS*.
- Liang Qu, Huaisheng Zhu, Ruiqi Zheng, Yuhui Shi, and Hongzhi Yin. 2021. Imgagn: Imbalanced network embedding via generative adversarial graph networks. In *KDD*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP-IJCNLP*.
- Min Shi, Yufei Tang, Xingquan Zhu, David Wilson, and Jianxun Liu. 2020. Multi-class imbalanced graph convolutional network learning. In *IJCAI*.
- Mingchen Sun, Kaixiong Zhou, Xin He, Ying Wang, and Xin Wang. 2022. Gppt: Graph pre-training and prompt tuning to generalize graph neural networks. In *KDD*.
- Xiangguo Sun, Hong Cheng, Jia Li, Bo Liu, and Jihong Guan. 2023. All in one: Multi-task prompting for graph neural networks. In *KDD*.
- Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. 2023. Does synthetic data generation of llms help clinical text mining? *arXiv preprint arXiv:2303.04360*.
- Llama team. 2024. The llama 3 herd of models. <https://ai.meta.com/research/publications/the-llama-3-herd-of-models/>.
- Junjiao Tian, Yen-Cheng Liu, Nathaniel Glaser, Yen-Chang Hsu, and Zsolt Kira. 2020. Posterior recalibration for imbalanced datasets. In *NeurIPS*.
- Yijun Tian, Kaiwen Dong, Chunhui Zhang, Chuxu Zhang, and Nitesh V Chawla. 2023. Heterogeneous graph masked autoencoders. In *AAAI*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *ICLR*.
- Xiao Wang, Deyu Bo, Chuan Shi, Shaohua Fan, Yanfang Ye, and S Yu Philip. 2022. A survey on heterogeneous graph embedding: methods, techniques, applications and sources. *IEEE Transactions on Big Data*.

- Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *WWW*.
- Xiao Wang, Nian Liu, Hui Han, and Chuan Shi. 2021. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *KDD*.
- Zehong Wang, Qi Li, Donghua Yu, Xiaolong Han, Xiao-Zhi Gao, and Shigen Shen. 2023. Heterogeneous graph contrastive multi-view learning. In *SDM*.
- Zehong Wang, Sidney Liu, Zheyuan Zhang, Tianyi Ma, Chuxu Zhang, and Yanfang Ye. 2025a. Can llms convert graphs to text-attributed graphs? In *NA*.
- Zehong Wang, Zheyuan Liu, Tianyi Ma, Jiazheng Li, Zheyuan Zhang, Xingbo Fu, Yiyang Li, Zhengqing Yuan, Wei Song, Yijun Ma, et al. 2025b. Graph foundation models: A comprehensive survey. *arXiv preprint arXiv:2505.15116*.
- Zehong Wang, Zheyuan Zhang, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2024. GFT: Graph foundation model with transferable tree vocabulary. In *NeurIPS*.
- Zehong Wang, Zheyuan Zhang, Tianyi Ma, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2025c. Neural graph pattern machine. In *ICML*.
- Zehong Wang, Zheyuan Zhang, Tianyi Ma, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2025d. Towards graph foundation models: Learning generalities across graphs via task-trees. In *ICML*.
- Zehong Wang, Zheyuan Zhang, Tianyi Ma, Chuxu Zhang, and Yanfang Ye. 2025e. Scalable graph generative modeling via substructure sequences. *arXiv preprint arXiv:2505.16130*.
- Zehong Wang, Zheyuan Zhang, Chuxu Zhang, and Yanfang Ye. 2025f. Training mlps on graphs without supervision. In *WSDM*.
- Qianlong Wen, Zhongyu Ouyang, Chunhui Zhang, Yiyue Qian, Chuxu Zhang, and Yanfang Ye. 2024. Gcvr: reconstruction from cross-view enable sufficient and robust graph contrastive learning. In *UAI*.
- Qianlong Wen, Zhongyu Ouyang, Jianfei Zhang, Yiyue Qian, Yanfang Ye, and Chuxu Zhang. 2022. Disentangled dynamic heterogeneous graph learning for opioid overdose prediction. In *KDD*.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks. In *ICLR*.
- Yanfang Ye, Yujie Fan, Shifu Hou, Yiming Zhang, Yiyue Qian, Shiyu Sun, Qian Peng, Mingxuan Ju, Wei Song, and Kenneth Loparo. 2020a. Community mitigation: A data-driven system for covid-19 risk assessment in a hierarchical manner. In *CIKM*.
- Yanfang Ye, Shifu Hou, Yujie Fan, Yiming Zhang, Yiyue Qian, Shiyu Sun, Qian Peng, Mingxuan Ju, Wei Song, and Kenneth Loparo. 2020b. α -satellite: An ai-driven system and benchmark datasets for dynamic covid-19 risk assessment in the united states. *IEEE Journal of Biomedical and Health Informatics*.
- Minji Yoon, John Palowitch, Dustin Zelle, Ziniu Hu, Ruslan Salakhutdinov, and Bryan Perozzi. 2022. Zero-shot transfer learning within a heterogeneous graph via knowledge transfer networks. In *NeurIPS*.
- Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. In *NeurIPS*.
- Xingtong Yu, Yuan Fang, Zemin Liu, and Xinming Zhang. 2024. Hgprompt: Bridging homogeneous and heterogeneous graphs for few-shot prompt learning. In *AAAI*.
- Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla. 2019. Heterogeneous graph neural network. In *KDD*.
- Yiming Zhang, Yiyue Qian, Yujie Fan, Yanfang Ye, Xin Li, Qi Xiong, and Fudong Shao. 2020. dstyle-gan: Generative adversarial network based on writing and photography styles for drug identification in darknet markets. In *ACSAC*.
- Yiming Zhang, Yiyue Qian, Yanfang Ye, and Chuxu Zhang. 2022. Adapting distilled knowledge for few-shot relation reasoning over knowledge graphs. In *SDM*.
- Zheyuan Zhang, Yiyang Li, Nhi Ha Lan Le, Zehong Wang, Tianyi Ma, Vincent Galassi, Keerthiram Murugesan, Nuno Moniz, Werner Geyer, Nitesh V Chawla, et al. 2025a. Ngqa: A nutritional graph question answering benchmark for personalized health-aware nutritional reasoning. In *ACL*.
- Zheyuan Zhang, Zehong Wang, Shifu Hou, Evan Hall, Landon Bachman, Jasmine White, Vincent Galassi, Nitesh V Chawla, Chuxu Zhang, and Yanfang Ye. 2024. Diet-odin: A novel framework for opioid misuse detection with interpretable dietary patterns. In *KDD*.
- Zheyuan Zhang, Zehong Wang, Tianyi Ma, Varun Sameer Taneja, Sofia Nelson, Nhi Ha Lan Le, Keerthiram Murugesan, Mingxuan Ju, Nitesh V Chawla, Chuxu Zhang, et al. 2025b. Mopi-hfrs: A multi-objective personalized health-aware food recommendation system with llm-enhanced interpretation. In *KDD*.
- Mingxia Zhao and Adele Lu Jia. 2024. Dahgn: Degree-aware heterogeneous graph neural network. *Knowledge-Based Systems*.
- Tianxiang Zhao, Xiang Zhang, and Suhang Wang. 2021. Graphsmote: Imbalanced node classification on graphs with graph neural networks. In *WSDM*.

Technical Appendix

A Detailed Related Works

Heterogeneous Graph Neural Networks. In the past few years, many heterogeneous graph neural networks (Qian et al., 2022b; Wang et al., 2022; Fu et al., 2020; Bing et al., 2023; Qian et al., 2023) have been proposed to learn node representations in heterogeneous graphs (Zhao and Jia, 2024; Zhang et al., 2024, 2025b; Wang et al., 2023; Ye et al., 2020b; Zhang et al., 2022; Qian et al., 2021a). For example, HAN (Fu et al., 2020) converts a heterogeneous graph into multiple metapath-based homogeneous graphs and leverages hierarchical attention to encode node-level and semantic-level structures. HetGNN (Zhang et al., 2019) introduces random walk (Perozzi et al., 2014) to sample node neighbors and further fuses their features using LSTMs (Hochreiter and Schmidhuber, 1997). HGT (Hu et al., 2020) proposes a transformer-based architecture for web-scale heterogeneous graphs. However, these models require sufficiently high-quality labeled data for training (Jiang et al., 2021; Yoon et al., 2022), which is labor-intensive and impractical for online drug trafficking detection. Unlike existing works (Qian et al., 2022a; Ye et al., 2020a; Qian et al., 2022c), we propose a graph prompt learning framework to address the label scarcity issue for drug trafficking detection on social media.

Prompt-based Learning on Graph. Recently, a new paradigm known as "pre-train and prompt" (Liu et al., 2023a) has gradually gained attention in the natural language processing field. This paradigm reformulates downstream tasks to resemble pre-train tasks via prompts among input texts (Gu et al., 2021). This not only bridges the gap between pretraining and downstream tasks but also instigates further research integrating prompting with pre-trained graph neural networks (Sun et al., 2023; Wang et al., 2025e; Ma et al., 2024; Wang et al., 2025c). For example, GraphPrompt (Liu et al., 2023c) utilizes the subgraph similarity and a learnable prompt for downstream few-shot learning. GPPT (Sun et al., 2022) introduces a template to align the pretext task of link prediction with downstream classifications. HGPrompt (Yu et al., 2024) proposes a dual-template design for heterogeneous graphs. However, the aforementioned models overlook the class imbalance issue in (heterogeneous) graphs, which

Algorithm 1: Training Procedure of LLM-HetGDT.

```
1 Input: Heterogeneous graph
    $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{A}, \mathcal{R})$  with plaintext set  $\mathcal{T}$ ,
   encoder  $f(\cdot)$  with parameters  $\phi$ , LLM
    $M(\cdot)$ 
2 # Pre-train Stage
3 while not converge do
4   Feed  $\mathcal{G}$  into encoder  $f$  to generate node
   embeddings  $z_i^{mp}$  and  $z_i^{sc}$ .
5   Compute the pretext loss  $\mathcal{L}_{cl}$  through
   Eq. 1. Optimize parameters  $\phi$  by
   minimizing the loss  $\mathcal{L}_{cl}$ .
6 # Fine-tune Stage
7 Freeze the parameters of encoder  $f$ .
8 Feed  $\mathcal{G}$  into LLM  $M$  to obtain the
   augmented heterogeneous graph  $\mathcal{G}'$ 
9 Initialize drug trafficking and node prompts,
    $\mathbf{C}$  and  $\mathbf{F}$ .
10 for each epoch do
11   Obtain prompted node features by
   applying feature prompt to node
   attribute features in  $\mathcal{G}'$ .
12   Feed the prompted node features into
   HGNNs  $f_{\phi^*}$  to generate target node
   embeddings  $\mathbf{Z}$ .
13   Compute the structure prompt  $\mathbf{S}$  via
   Eq. 8 - Eq. 21.
14   Update target node embeddings by
   applying the structure prompt.
15   Compute the total loss  $\mathcal{L}_{pt}$  via Eq. 13.
16   Optimize the parameters by minimizing
    $\mathcal{L}_{pt}$ .
```

is a common scenario in social media for drug trafficking activities.

B Data Description

B.1 HG Construction on Social Media Data

To comprehensively describe the drug trafficking activities on social media, we construct a drug trafficking HG that integrates informative content features and complex relationships among online users. In this paper, we regard Twitter as an example to study online drug trafficking activities. Specifically, we build an HG, with three types of nodes (user, tweet, and keyword) and six types of relations as well as features of each node, such that both content and relation information can be

Content Feature	
User (U)	username, user ID, profile, and # of followers.
Tweet (T)	text content, # of likes, reply, and retweet.
Keyword (K)	keywords extracted from text contents.
Relation	
R1: User-follow-User	R2: User-post-Tweet
R3: User-engage-Tweet	R4: User-profile-Keyword
R5: Tweet-include-Keyword	R6: Tweet-tag-Keyword

Table 5: Content and relation information in heterogeneous graph dataset Twitter-HetDrug.

exploited simultaneously. The details of content features and relationship information are discussed below.

Content Feature. To comprehensively characterize each target node (user node), unlike existing works (Liu et al., 2021; Ma et al., 2023) that merely consider users as nodes, we introduce three types of nodes, i.e., User (U), Tweet (T), and Keyword (K), into HG. Specifically, for each node $v_i^A \in \mathcal{V}_A$, we concatenate the text content (as listed in Table 5), denoted as t_i^A , and further leverage the pre-trained transformer-based language model, SentenceBert (Reimers and Gurevych, 2019), to convert the concatenated text information into a fixed-length feature vector as attribute feature x_i^A . Here, x_i^A denotes the attribute feature of node v_i^A belongs to node set \mathcal{V}_A with node type $A \in \mathcal{A}$, and $\mathcal{A} = \{U, T, K\}$.

Relations. To determine whether a user is a drug trafficking participant on social media, we consider not only the content-based features but also the complex relationships among users, tweets, and keywords. Particularly, we define six kinds of relationships (R1-R6 in Table 5): *R1: user-follow-user* denotes a user is following or followed by another user; *R2: user-post-tweet* represents a user posts a tweet; *R3: user-engage-tweet* denotes that a user engages into a tweet, including likes, replies, retweet, and mention; *R4: user-profile-keyword* represents that a user’s profile contains the keyword; *R5: tweet-include-keyword* denotes that a tweet content includes a keyword in plaintext; *R6: tweet-tag-keyword* indicates that a tweet tags a keyword as a hashtag. To summarize, we build a HG $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X}, \mathcal{A}, \mathcal{R})$ by integrating content features of three kinds of entities and six types of relationships among these entities. The data statistic of the constructed HG is listed in Table 6.

B.2 Data Collection on Social Media

We crawl metadata through the official Twitter API from Dec. 2020 to Aug. 2021. Afterward, following the existing work (Ma et al., 2023; Qian et al., 2021b), we generate a keyword list that covers 21 drug types that may cause drug overdose or drug addiction problems to filter the tweets that contain drug-relevant information. The drug types and partial drug keywords are listed in Appendix B.4. Based on the keyword list, we obtain 54,680 users, with corresponding 326,826 tweets and 288 valid keywords.

B.3 Annotation Rules

With the filtered metadata, six researchers spend 62 days annotating these Twitter users separately, following the annotation rules discussed below: (i) If a user actively promotes some type of drug on Twitter or has rich connections (e.g., following, replying, liking, and retweeting) with other drug-related users in specific drug communities, he/she will be considered a drug trafficking participant. (ii) If a user appears to suffer from drug overdose or drug addiction to a specific drug, he/she will be identified as a drug trafficking participant. (iii) If we can find evidence on Twitter that a user used to purchase specific drugs from others on Twitter, then he/she is considered a drug trafficking participant. To ensure fairness and mitigate bias, for these Twitter users with disagreed labels, we conducted further discussion among annotators for cross-validation. Based on the above strategy, we can obtain the ground truth for the drug trafficking detection task.

B.4 Drug Types and Keywords

Following previous works (Qian et al., 2021b; Ma et al., 2023; Zhang et al., 2025a), we generate a keyword list that covers 21 drug types that may cause drug overdoses or drug addiction problems

to filter the tweets that contain drug-relevant information. Table 7 lists the drug types and partial drug keywords.

Type	Category	Count
Node	User	28,839
	Tweet	118,697
	Keyword	288
Class	Participant	3,437
	Benign	25,402
Edge	R1	100,571
	R2	118,697
	R3	44,197
	R4	64,997
	R5	260,868
	R6	39,129

Table 6: Statistics of the constructed HG. The underlined node type is the target node for classification.

Drug Type	Drug Name (Keywords)
Cannabis	Cannabis, and infused products
Opioid	Oxycodone, Codeine, etc.
Hallucinogen	LSD, MDMA, Shroom, DMT, etc.
Stimulant	Cocaine, Amphetamine, etc.
Depressant	Xanax, Valium, Halcion, etc.

Table 7: Drug types and partial drug keywords.

C Prompt Templates

In this section, we provide the node prompt template in Table 8 and edge prompt template in Table 9.

D Additional Mechanisms Details

D.1 Attention Mechanism

The attention mechanism $\text{Att}(\cdot)$ in Equation 6 is defined as:

$$\text{Att}(\mathbf{X}, \mathbf{F}) = \sum_{k=1}^K \frac{\exp(\sigma((\mathbf{f}_k)^\top \cdot \mathbf{X}_{i,:}))}{\sum_{j=1}^K \exp(\sigma((\mathbf{f}_j)^\top \cdot \mathbf{X}_{i,:}))} \cdot \mathbf{f}_k, \quad (14)$$

where σ denotes the activation function, and \mathbf{f}_k is the k -th row of the feature token \mathbf{F} .

D.2 Set Transformer

Here, we present set transformer in Equation 6. Particularly, the set transformer $\text{PMA}(\cdot)$ is defined as:

$$\text{PMA}(\mathcal{S}) = \text{LN}(\mathbf{Q} + \text{MLP}(\mathbf{Q})), \quad (15)$$

$$\mathbf{Q} = \text{LN}(\psi + \text{MH}(\psi, \mathbf{Y}, \mathbf{Y})), \quad (16)$$

INPUT:

Instruction: Act as an expert in social network analysis, your task is to generate a synthetic user based on the given context, including username, user ID, and user profile. The generated synthetic user information should capture the **key information** in the original user while **ensuring variability**.

Context: Username: **<Username>**; User ID: **<User ID>**; User Profile: **<User Profile>**.

Table 8: The prompt template in LLM-HetGDT to generate synthetic online drug trafficking user information.

INPUT:

Instruction: Act as an expert in social network analysis; given the original user information, the connected neighbor information with neighbor type, and synthetic user information generated from the original user, your task is to determine which neighbors the synthetic user should connect to.

Context: Original User Information: **<Original User Information>**; Neighbor Type: **<Neighbor Type>**; Neighbor Information: [**<Neighbor Information 1>**, ...]; Synthetic User Information: **<Synthetic User Information>**.

Table 9: The prompt template in LLM-HetGDT to generate neighbors of synthetic online drug trafficking users.

where \mathcal{S} is the set of input tokens, $\text{LN}(\cdot)$ denotes the layer normalization, $\text{MLP}(\cdot)$ is multi-layer perceptron, $\psi \in \mathbb{R}^{d_T}$ is a learnable parameters, served as a seed (query) vector, $\mathbf{Y} \in \mathbb{R}^{|\mathcal{C}| \times d_T}$ is the stacked embeddings of set \mathcal{S} , and $\text{MH}(\cdot)$ represents the multi-head attention mechanism. Besides, the multi-head attention mechanism $\text{MH}(\cdot)$ is formulated as:

$$\text{MH}(\psi, \mathbf{Y}, \mathbf{Y}) = \|\|_{j=1}^h \omega(\psi_j \cdot \mathbf{K}_j^\top) \mathbf{V}_j, \quad (17)$$

$$\mathbf{K}_j = \mathbf{Y} W_j^K, \mathbf{V}_j = \mathbf{Y} W_j^V, \quad (18)$$

where $\|\|$ denotes the concatenation operation, h is the number of heads, ω is the softmax function, and $\psi_j \in \mathbb{R}^{d_T^M}$ with dimension $d_T^M = d_T/h$. Here $\psi = \|\|_{j=1}^h \psi_j$, $W_j^K \in \mathbb{R}^{d_T \times d_T^M}$ and $W_j^V \in \mathbb{R}^{d_T \times d_T^M}$ are the weight matrices for key and value vectors, respectively.

D.3 Meta Path

A meta-path P is a path defined in the form of $A_1 \xrightarrow{R_1} A_2 \xrightarrow{R_2} \dots \xrightarrow{R_l} A_{l+1}$, where $A_i \in \mathcal{A}$ and $R_i \in \mathcal{R}$. It describes a composite relation

$R = R_1 \circ R_2 \circ \dots \circ R_l$ between nodes A_1 and A_{l+1} . Here \circ denotes the composition operator. In Twitter-HetDrug, we define the following three types of meta paths:

- User-Tweet-User.
- User-Tweet-Keyword-Tweet-User.
- User-Keyword-User.

D.4 Semantic Attention Mechanisms

The semantic attention mechanism **SM-Att**(\cdot) in Equation 8 is defined as:

$$\mathbf{SM-Att}(\mathcal{H}) = \sum_{\mathbf{H}_i \in \mathcal{H}} \alpha_i \cdot \mathbf{H}_i, \text{ where} \quad (19)$$

$$\alpha_i = \frac{\exp(w_i)}{\sum_{H_j \in \mathcal{H}} \exp(w_j)}, \text{ and} \quad (20)$$

$$w_i = \frac{1}{|\mathcal{H}|} \sum_{\mathbf{H}_i \in \mathcal{H}} q^\top \cdot \tanh(W \cdot \mathbf{H}_i + b), \quad (21)$$

where $q \in \mathbb{R}^d$ is the semantic-level attention vector, $W \in \mathbb{R}^{d \times d}$ denotes the weight matrix, and $b \in \mathbb{R}^d$ is the bias vector. The parameters W and b are shared across all meta-paths.

E Baseline Settings

We compare LLM-HetGDT with seven groups of baseline methods, including graph representation learning methods (G1), heterogeneous graph representation learning models (G2), heterogeneous graph self-supervised learning methods (G3), generic methods against class imbalance data (G4), graph models against the class imbalance issue (G5), graph prompt learning models (G6), and LLMs-based methods (G7). Next, we would like to introduce these baseline methods in detail.

G1 Graph Representation Learning Methods:

We implement Graph Convolution Network (GCN) (Kipf and Welling, 2017), Graph Attention Network (GAT) (Veličković et al., 2018), and Graph Isomorphism Network (GIN) (Xu et al., 2019). We employ a two-layer MLP as a classifier at the end of these models for the node classification tasks.

G2 Heterogeneous Graph Representation Learning Models:

We choose three state-of-the-art (SOTA) heterogeneous graph representation learning methods, including Heterogeneous Graph Neural Network (HetGNN) (Zhang et al.,

2019), Heterogeneous Graph Attention Network (HAN) (Wang et al., 2019), and Heterogeneous Graph Transformer (HGT) (Hu et al., 2020). Specifically, HetGNN utilizes random walks to sample node neighbors and further fuse their attribute features using LSTMs to generate node embeddings. HAN converts multiple metapath-based homogeneous graphs and leverages hierarchical attention to encode node-level and semantic-level structures to obtain node embeddings. HGT adopts a transformer-based architecture for web-scale heterogeneous graphs. Similarly, we utilize two-layer MLP as a classifier over the node embeddings for the node classification tasks.

G3 Heterogeneous Graph Self-supervised Learning Methods:

We also implement four self-supervised methods for heterogeneous graph representation learning, i.e., Deep Graph Infomax for Attribute Multiplex Network (DMGI) (Park et al., 2020), Heterogeneous Graph Neural Network with Co-contrastive Learning (HeCo) (Wang et al., 2021), Heterogeneous Graph Masked Autoencoder (HGMAE) (Tian et al., 2023), and MetaHG (Qian et al., 2021b). Particularly, DMGI introduces a consensus regularization to minimize the disagreements among the relation-type specific node embeddings and a universal discriminator that discriminates true samples regardless of the relation types. HeCo adopts cross-view contrastive learning on heterogeneous graphs, i.e., meta-path and network schema view encoders. HGMAE leverages two masking strategies toward meta-path adjacency matrix and attributes features with three training techniques, i.e., meta-path edge reconstruction, target attribute restoration, and positional feature prediction, to obtain target node embeddings. MetaHG is a meta knowledge distillation framework over heterogeneous graphs for drug trafficking detection.

G4: Generic Methods Against Class Imbalance Data:

We compare LLM-HetGDT with three classic generic methods for class imbalance data. Re-weighting (Cui et al., 2019) scales the training loss with respect to the weight of classes. Over-sampling (Kang et al., 2020) duplicates the samples from minority classes to balance the training data distribution. SMOTE (Chawla et al., 2002) is an over-sampling method that generates synthetic samples from minority classes by interpolating between existing samples.

G5 Graph Models Against the Class Im-

Setting		Twitter-HetDrug 10%	Twitter-HetDrug 20%	Twitter-HetDrug 40%			
Group	Model	Macro-F1	GMean	Macro-F1	GMean	Macro-F1	GMean
G1	GCN	55.74 \pm 3.71	53.92 \pm 2.45	58.38 \pm 2.39	56.43 \pm 2.46	60.94 \pm 2.33	59.72 \pm 2.07
	GAT	54.24 \pm 3.53	52.06 \pm 2.52	57.93 \pm 3.38	55.04 \pm 2.74	60.63 \pm 2.80	59.30 \pm 2.75
	GIN	56.92 \pm 1.14	54.26 \pm 2.37	59.15 \pm 1.61	57.92 \pm 2.37	61.09 \pm 3.74	60.67 \pm 2.83
G2	HetGNN	58.52 \pm 3.56	59.21 \pm 2.36	60.86 \pm 2.71	61.39 \pm 2.96	63.46 \pm 3.01	64.05 \pm 3.13
	HAN	60.25 \pm 0.79	59.26 \pm 2.72	63.58 \pm 0.95	62.07 \pm 2.47	66.71 \pm 0.92	65.98 \pm 3.04
	HGT	60.95 \pm 1.61	59.88 \pm 2.09	63.85 \pm 1.95	63.70 \pm 2.59	67.08 \pm 2.06	66.37 \pm 3.84
G3	DMGI	60.53 \pm 3.62	59.04 \pm 2.65	63.26 \pm 3.19	62.94 \pm 2.90	65.28 \pm 2.57	64.95 \pm 3.05
	HeCo	62.93 \pm 2.02	60.85 \pm 1.05	64.83 \pm 2.53	63.21 \pm 3.62	67.94 \pm 2.92	66.52 \pm 2.69
	HGMAE	61.84 \pm 2.61	60.01 \pm 2.54	62.53 \pm 3.01	61.73 \pm 3.61	65.74 \pm 2.46	64.63 \pm 2.90
	MetaHG	61.77 \pm 2.45	60.35 \pm 2.98	62.16 \pm 3.58	62.50 \pm 3.21	66.41 \pm 2.07	65.39 \pm 2.47
G4	Re-weighting	60.72 \pm 0.89	60.31 \pm 0.94	62.38 \pm 0.52	63.07 \pm 0.58	65.05 \pm 0.59	65.91 \pm 0.46
	Over-sampling	63.79 \pm 0.73	62.27 \pm 0.75	65.10 \pm 0.32	63.08 \pm 0.38	68.35 \pm 0.40	67.39 \pm 0.85
	SMOTE	64.25 \pm 0.72	62.17 \pm 0.86	66.94 \pm 0.82	63.62 \pm 0.77	69.24 \pm 0.84	66.77 \pm 0.93
G5	GraphENS	65.37 \pm 1.72	63.80 \pm 3.49	66.93 \pm 0.94	64.06 \pm 3.19	68.37 \pm 1.37	70.82 \pm 2.29
	GraphSMOTE	65.82 \pm 3.75	64.04 \pm 2.39	68.07 \pm 3.89	65.83 \pm 2.63	70.32 \pm 2.46	69.25 \pm 2.09
	PC-GNN	65.98 \pm 2.83	63.84 \pm 3.48	67.52 \pm 2.95	64.71 \pm 3.04	70.59 \pm 2.48	71.94 \pm 3.74
G6	GPPT	62.84 \pm 4.72	60.41 \pm 2.40	64.97 \pm 3.47	63.28 \pm 2.53	68.04 \pm 2.46	67.42 \pm 2.73
	Self-Pro	63.37 \pm 3.52	61.05 \pm 3.61	65.23 \pm 3.14	63.94 \pm 3.51	68.33 \pm 2.73	67.16 \pm 1.93
	GraphPrompt	64.37 \pm 1.37	62.38 \pm 3.52	65.89 \pm 2.85	64.53 \pm 1.92	68.62 \pm 2.46	68.37 \pm 2.18
	HGPrompt	66.85 \pm 2.95	64.37 \pm 2.53	67.97 \pm 3.52	66.08 \pm 2.08	69.72 \pm 3.92	69.42 \pm 3.24
G7	LLama 70B	65.07 \pm 0.00	63.84 \pm 0.00	64.90 \pm 0.00	63.77 \pm 0.00	75.37 \pm 0.00	64.02 \pm 0.00
	Mistral 24B	64.83 \pm 0.00	62.87 \pm 0.00	64.90 \pm 0.00	62.57 \pm 0.00	64.73 \pm 0.00	62.36 \pm 0.00
	DeepSeek-R1 70B	65.74 \pm 0.00	63.92 \pm 0.00	65.36 \pm 0.00	63.98 \pm 0.00	65.18 \pm 0.00	63.65 \pm 0.00
Ours	LLM-HetGDT	70.52 \pm 1.58	71.69 \pm 2.50	71.36 \pm 1.72	73.87 \pm 1.45	74.06 \pm 2.74	76.24 \pm 2.85

Table 10: Performance (mean \pm std) comparison among various methods on Twitter-HetDrug for drug trafficking detection. Bolded numbers indicate the best performance, and underline numbers represent runner-up performance.

Hyper-parameters	Values
Hidden Dimension	{128, 256, 512, 1024}
Number of Layers	{1, 2, 3}
Normalize	{none, batch, layer}
Learning Rate	{1e-2, 1e-3, 1e-4}
Weight Decay	{0.0, 1e-5, 1e-4, 1e-3}
Dropout	{0.0, 0.1, 0.5, 0.8}
Attention Head*	{1, 2, 4, 8}

Table 11: The hyper-parameters we searched. Star hyper-parameters (*) are not applicable to every models.

balance Issue: We select three SOTA graph-based methods for class imbalance issues, including GraphENS (Park et al., 2022), GraphSMOTE (Zhao et al., 2021), and PC-GNN (Liu et al., 2021). GraphENS synthesizes the whole ego network for the minor class (minor node and its one-hop neighbors) by combining two different ego networks based on their similarity and further introduces a saliency-based node mixing method to exploit the abundant class-generic attributes of other nodes while blocking the injection of class-specific features. GraphSMOTE extends SMOTE to the graph data by designing a re-sampling module and an edge predictor to

determine the connectivity between synthetic nodes and other nodes in graphs. PC-GNN devises a pick-and-choose strategy to handle the class imbalance issue in multi-relation graphs.

G6 Graph Prompt Learning Models: We also compare our proposed method with four (heterogeneous) graph prompt learning baseline methods, i.e., GPPT (Sun et al., 2022), Self-Pro (Gong et al., 2024), GraphPrompt (Liu et al., 2023c), and HGPrompt (Yu et al., 2024). GPPT and GraphPrompt introduce prompt templates to align the pretext task of linked prediction with downstream classification. Self-Pro leverages asymmetric graph contrastive learning as a pretext to address heterophily and align the objective of pretext and downstream tasks. It also proposes to reuse the component from pre-training as an adapter and introduce self-prompts based on the graph for task adaptations. HGPrompt designs a dual prompt template for both homogeneous and heterogeneous graphs. Mention that for all baseline methods that are designed for homogeneous graphs, following existing works (Wang et al., 2019; Park et al., 2020; Tian et al., 2023), we directly use all meta-paths with target node attribute features as homogeneous

Ratio	Dim.	Layers	Normalize	LR.	Decay	Dropout	Att. Head	# basic vec.	δ	λ
10%	256	3	layer	1e-3	1e-4	0.1	4	10	5e-2	1e-3
20%	512	2	layer	1e-3	1e-5	0.5	1	15	5e-2	1e-3
40%	512	3	layer	1e-4	1e-5	0.5	1	10	5e-2	1e-3

Table 12: The hyper-parameters for our LLM-HetGDT in basic training setting.

graphs for downstream tasks. Additionally, we strictly follow the source code to reproduce the experiment result.

G7 LLMs-based Methods: As there is only one LLM-based method, i.e., KG-GPT (Hu et al., 2024), for drug trafficking detection, we employ the prompt in KG-GPT with different backbone LLMs as baseline methods. Specifically, we choose three state-of-the-art LLMs, including Llama-3.1-70B-Instruct (team, 2024), Mistral-Small-24B-Instruct (Hugging Face, Mistral AI, 2024), and DeepSeek-R1-Distill-Llama-70B (DeepSeek-AI, 2025).

F Additional Experimental Settings

We follow the existing works (Wang et al., 2025a; Qian et al., 2021b) to grid search the optimal hyperparameter of each model, and report the best performance. The hyperparameters we search for are listed in Table 11. The optimal hyperparameters of our model are provided in Table 12.

G Deployment Evaluation Settings

To further validate the applicability of our model, we employ LLM-HetGDT to detect drug trafficking activities in a live manner. Specifically, we first identify drug sellers in our dataset, i.e., Twitter-HetDrug, and further leverage these anchor users to expand their social circles over Twitter. Then we follow the same procedure in Appendix B to construct the HG for the expanded social circles and annotate the labels for the users. The constructed HG contains 416 users, including 61 drug trafficking participants and 355 benign users, 3,209 tweets, and 84 keywords. Afterward, we employ baseline methods and our proposed method to detect drug trafficking participants over the constructed HG. Mention that, we randomly choose three drug trafficking participants and seven benign users as the training set, and the rest are for the test set. To ensure a fair comparison between supervised methods and LLM-based approaches, we include the three drug trafficking participants from the training set as few-shot examples within the LLM prompt

during inference.

H Additional Experiment Comparison

Due to the space limit in the main paper, we provide additional experiment results in this section. Table 10 lists the performance of all baseline methods and our proposed method over Twitter-HetDrug.

# of basis vectors	1	5	10	15
Twitter 10%	66.83	<u>69.78</u>	70.52	69.32
Twitter 20%	68.39	70.69	<u>71.04</u>	71.36
Twitter 40%	71.24	<u>73.72</u>	74.06	73.58
Hidden dimension	128	256	512	1,024
Twitter 10%	69.74	70.52	<u>70.28</u>	69.84
Twitter 20%	70.36	<u>71.09</u>	71.36	70.48
Twitter 40%	73.85	<u>73.97</u>	74.06	73.89

Table 13: Sensitivity test (Macro-F1) on the hyper-parameters: # of basis vectors and hidden dimensions.

I Hyper-parameter Sensitivity Analysis

We conduct hyperparameter sensitivity experiments among Twitter-HetDrug. Specifically, we control the number of basis vectors in node prompts, i.e., K , and evaluate the performance in Macro-F1 of LLM-HetGDT. The performance is shown in Table 13. The bolded numbers indicate the best performance, and the underlined numbers represent the runner-up performance. According to Table 13, we find that even a relatively small number ($K = 5$) of basic vectors can achieve good performance in drug trafficking detection. The best performance is achieved at $K = 10$ for Twitter-HetDrug with train split 10% and 40%, and shows satisfactory performance for Twitter-HetDrug with train split 20%. This indicates the effectiveness of our model, as it achieves excellent performance with a relatively small number of basis vectors. We also perform a sensitivity analysis to show the impact of hidden dimensions. According to Table 13, we find that the optimal hidden dimension can be different across Twitter-HetDrug with different training splits. For example, the optimal hidden dimension for the 10% training split is 256,

while 512 works best for the 20% and 40% splits. However, increasing the hidden dimension to 1024 leads to reduced model performance, possibly because a wider hidden layer may prevent the model from focusing on the most relevant information. Hence, using standard hidden dimensions, i.e., 256 or 512, is sufficient for LLM-HetGDT to capture the complex information, which again shows the superiority of our model LLM-HetGDT.

J Additional Case Study

We further conduct a case study for edge generation. Specifically, after obtaining synthetic users, we employ LLM to generate edges between synthetic users and the neighbors of corresponding labeled users. Table 14 illustrates the edges generated for synthetic users in Section 4.6. According to Table 14, the original user has four neighboring tweets, consisting of one daily update and three advertising posts. In the response, the LLM accurately identifies the last three advertising posts as the neighbors for the synthetic user while disregarding the first. Moreover, it offers a logical explanation for each selection. For instance, the third tweet, "Same day shipping in ***, " lacks any drug-related content, which might be ignored by edge generation methods at the attribute feature level (Zhao et al., 2021; Liu et al., 2021). However, the LLM can accurately recognize this tweet as a promotion post for the same-day shipping service, effectively capturing the underlying behavior of the synthetic user.

INPUT:

Instruction:

Act as an expert in social network analysis; given the original user information, the connected neighbor information with neighbor type, and synthetic user information generated from the original user, your task is to determine which neighbors the synthetic user should connect to.

Please provide the reasons behind your selection.

Context:

Original User Information: Username: yas***exb; User ID: @dru***cy; User Profile: Online dispensary, premium lab-tested psychedelic products, must be 18+, same day shipping in ***. DM me for menu. #LSD #420 #acid.

Neighbor Type: Tweet.

Neighbor Information:

1. Text Content: Vibe for everyday. How was your day?
2. Text Content: DM me for menu. #Meth#LSD#420#AcidTrip.
3. Text Content: Same day shipping in ***.
4. Text Content: !!!Giveaway!!! I will be giving away ONE 12.30OZ BAG OF METH. To enter, - Follow me, Retweet, and Reply 'Meth' under this post. END AUG ***.

Synthetic User Information:

Username: Psy***ete; User ID: @ss***ete; User Profile: Your one-stop shop for lab-tested psychedelic products. Get ready to trip with the selection of premium psychedelic products. #lsdtrip #psychedelics #acid #weed #420 #mushrooms.

OUTPUT:

Selected Neighbors: [2, 3, 4].

Reasons:

Post 2 contains the hashtags #Meth, #LSD, #420, and #AcidTrip, and [these drug names match the products in the profile](#).

Post 3 [promotes the same-day shipping service in location ***](#), which is most likely a post by the synthetic user.

Post 4 is a giveaway post about meth. This post mainly focuses on [giveaway and promotion](#) the products to attract potential buyers, which is highly relevant to the synthetic user's profile.

Table 14: Example of neighbors selected by LLM.