Abnormal Event Detection via Hypergraph Contrastive Learning

Bo Yan * Cheng Yang * Chuan

Chuan Shi ^{*†} Jiawei Liu ^{*}

Xiaochen Wang *

Abstract

Abnormal event detection, which refers to mining unusual interactions among involved entities, plays an important role in many real applications. Previous works mostly oversimplify this task as detecting abnormal pair-wise interactions. However, real-world events may contain multi-typed attributed entities and complex interactions among them, which forms an Attributed Heterogeneous Information Network (AHIN). With the boom of social networks, abnormal event detection in AHIN has become an important, but seldom explored task. In this paper, we firstly study the unsupervised abnormal event detection problem in AHIN. The events are considered as star-schema instances of AHIN and are further modeled by hypergraphs. A novel hypergraph contrastive learning method, named AEHCL, is proposed to fully capture abnormal event patterns. AEHCL designs the intra-event and inter-event contrastive modules to exploit self-supervised AHIN information. The intra-event contrastive module captures the pair-wise and multivariate interaction anomalies within an event, and the inter-event module captures the contextual anomalies among events. These two modules collaboratively boost the performance of each other and improve the detection results. During the testing phase, a contrastive learning-based abnormal event score function is further proposed to measure the abnormality degree of events. Extensive experiments on three datasets in different scenarios demonstrate the effectiveness of AE-HCL, and the results improve state-of-the-art baselines up to 12.0% in Average Precision (AP) and 4.6% in Area Under Curve (AUC) respectively.

1 Introduction

Events widely exist in social networks, which consist of several entities and complex interactions among them. Abnormal events refer to unusually or rarely happened events among observed event samples. Figure 1(d) shows a toy example of abnormal events in the citation network, which is considered abnormal since the involved entities exhibit rare interactions (data-mining specialists seldom collaborate on a COVID-19 paper



Figure 1: A toy example of abnormal events in AHIN. (a) A citation network. (b) Network schema of the citation network. (c) A star schema instance (event) guided by network schema. (d) Events modeled by hypergraph.

with radiologists). To date, abnormal event detection has become a significant task in fraud detection [26], Internet of things (IoT) security [25], computer network monitoring [4], and other real applications, thus drawing increasing concerns in recent years.

Existing approaches mainly focus on detecting events with pair-wise abnormal patterns [7, 4, 1]. Moreover, some studied events only have simple structures [22, 26] or categorical attributes [25, 9]. However, realworld events may contain multiple types of entities with rich attributes, as well as complex interactions among them, which forms an Attributed Heterogeneous Information Network (AHIN) [12]. As depicted in Figure 1(a), the event of publishing a paper in AHIN is associated with a set of multi-typed attributed nodes, making the event not limited to simple pair-wise interactions (e.g., a data mining specialist writes a COVID-19 paper), but complex set-wise interactions (e.g., several data mining specialists cooperate to write a COVID-19 paper on a bioinformatics venue). Detecting abnormal events in AHIN became an essential, but largely neglected task in social network analysis and abnormal

^{*}Beijing University of Posts and Telecommunications, China. Email: {boyan, yangcheng, shichuan, liu_jiawei, 2018211801}@bupt.edu.cn

[†]Corresponding author

detection. For example, detecting abnormal collaborations in the academic network can better understand research trends and promote cooperation.

Although this task is highly urged to be explored, it's never a trivial task. First, the complex set-wise patterns in AHIN make the modeling of abnormal events more challenging. Traditional pair-wise modeling loses its effects since all pair-wise interactions may be normal but the set-wise interaction is abnormal. In this regard, a general framework should be proposed to model events in AHIN, as well as fully consider complex abnormal setwise patterns. Second, due to the scarcity of anomalies and the prohibitive cost of labeling processes, we need to conduct abnormal event detection in an unsupervised manner, which compounds the difficulties of capturing abnormal patterns in AHIN. Finally, to measure the abnormal degrees in an unsupervised setting, a proper abnormal event score function is desired, which is capable of truly reflecting the abnormal degrees.

To tackle these challenges, in this paper, we study the problems of unsupervised abnormal event detection in AHIN and propose a novel framework for Abnormal $\underline{\mathbf{E}}$ vent detection via $\underline{\mathbf{H}}$ ypergraph $\underline{\mathbf{C}}$ ontrastive $\underline{\mathbf{L}}$ earning (AEHCL). First, we consider an event in AHIN as a star-schema instance [18] and further use the hyperedge [8, 19] to model the set-wise interactions. A hyperedge in the hypergraph is associated with more than two entities, endowing the entities within an event not limited to pair-wise interactions. Second, based on this event modeling, we propose a hypergraph contrastive learning method to fully capture abnormal patterns in an unsupervised manner, which includes two contrastive strategies at the intra-event and inter-event levels. At the intra-event level, the pair-wise contrastive module captures abnormal interactions between node pairs within an event, and the multivariate contrastive module further considers high-order interactions. At the inter-event level, we contrast events with their neighbors to capture the inconsistent patterns between local events. All these modules are optimized in an endto-end manner. Finally, to measure the abnormal degrees of events in AHIN, we further propose a novel contrastive-based abnormal event score function, which effectively integrates the detection results of the abovementioned modules.

The contributions are summarized as follows:

(1) To the best of our knowledge, this is the first attempt to study the problem of abnormal event detection in AHIN, which is an important and practical task in many application scenarios.

(2) We propose a novel hypergraph contrastive learning method, called AEHCL, leveraging both intraand inter-event contrastive tasks to capture abnormal event patterns within an event and between events. A contrastive-based abnormal event score function is further designed to measure the abnormal degrees.

(3) We conduct extensive experiments on three types of AHIN datasets. The results improve state-of-the-art baselines up to 12.0% in AP and 4.6% in AUC scores respectively, which demonstrate the effectiveness of our model.

2 Related Work

 $\mathbf{2.1}$ Abnormal Event Detection The abnormal event usually implies that a set of entities as well as their interactions rarely co-occurrence. For a set of entities with abnormal structures, [22] detects suspicious dense blocks by utilizing an anomaly-aware representation method. [26] further maps different types of nodes into a shared latent space. These models haven't formally defined abnormal events and also ignore rich attributes of entities. Some other works study discrete event sequence abnormal detection [3, 27, 10], where each event is represented by a categorical value [25, 9]. The event in the HIN is first defined by [8] as a node set that forms a complete semantic unit. [4] proposes APE to model the weighted pair-wise interactions of the heterogeneous categorical event [1]. AEHE [7] further incorporates node attributes and node's secondorder representations to detect abnormal events with specific meta-path. Besides these abnormal event detection studies, the most relevant task is known as abnormal node detection. CoLA [15], and ANEMONE [14] are two advanced works using contrastive learning to model mismatching patterns between neighbor nodes. Both these works only model abnormal nodes or abnormal events with simple patterns but fail to capture complex abnormal interactions of events in AHIN.

2.2 Heterogeneous Information Network Analysis Heterogeneous information networks (HIN) [18] analysis focuses on modeling complex entities and their rich relations in various applications. They mainly leverage meta-path to learn semantic-preserving representations of HIN. Metapath2vec [5] designs a metapath based random walk and utilizes skip-gram to learn HIN embedding. HAN [23] proposes node and semanticlevel attention to learn the importance of meta-paths. HeCo [24] introduces a cross-view contrastive mechanism to HIN analysis. HeteHG-VAE [6] further uses hypergraphs to model high-order complex interactions of HIN and proposes a hypergraph variational autoencoder to learn robust node representations. HIN analvsis has been widely used in fraud detection [12], recommendation [16], and other fields. However, few works use HIN analysis in abnormal event detection.

Contrastive Learning Contrastive learning has $\mathbf{2.3}$ achieved great success in representation learning [3, 11], the core idea of which is to contrast positive pairs against negative pairs. As for graph-related contrastive learning methods, DGI [21] proposes an unsupervised learning objective based on mutual information theory to contrast, MVGRL [10] maximizes the mutual information between the node representation of one view and the graph representation of another view. In the heterogeneous graph domain, DMGI [17] introduces a consistency regularisation framework that minimized divergence between the specific relation type of node embeddings. Besides, some works apply contrastive learning to abnormal node detection [15, 27]. However, there is no research on using contrastive learning for abnormal event detection in AHIN.

3 Preliminary

In this paper, we perform abnormal event detection in the Attributed Heterogeneous Information Network (AHIN), which can be formally defined as follows:

Definition 1. Attributed Heterogeneous Information Network (AHIN). An AHIN is denoted as G = (V, E, X) consisting of an object set V, a link set E and an attribute matrix $X \in \mathbb{R}^{|V| \times k}$. It is also associated with a object type mapping function $\phi : V \to \mathcal{A}$ and a link type mapping function $\varphi : E \to \mathcal{R}$. \mathcal{A} and \mathcal{R} denote the sets of predefined object and link types, where $|\mathcal{A}| + |\mathcal{R}| > 2$. The **network schema** [18] $S^G = (\mathcal{A}, \mathcal{R})$ can be seen as a meta template of an AHIN G, which is a graph defined over object types \mathcal{A} , with links as relations from \mathcal{R} .

Definition 2. Event in AHIN. An event $Q = (v^q, V^q, E^q, X^q, S^G)$ in an AHIN is an instance of star schema S^G , which includes an object set $V^q \subseteq V$, a link set $E^q \subseteq E$ and an attribute matrix $X^q \in R^{|V^q| \times k}$. $v^q \in V^q$ is the center node which uniquely identifies an event Q and the context node set $V^c = V^q \setminus v^q$ denotes the context of an event.

A toy example of AHIN and its network schema is shown in Figure 1(a) and (b) respectively for the citation network. Particularly, star schema network is a kind of AHIN, in which links only exist between objects with a center type and others. Guided by star schema, we can extract star schema instances (events) from the AHIN. Figure 1(c) shows an example of paper publication event. The center node is a paper uniquely identifying the event, and the context nodes are the venue and authors. Instead of modeling each event as a set of binary relations (links), we employ hypergraph [2] to model events by viewing involved nodes as a whole.

Definition 3. Hypergraph Modeling for Event.

A hypergraph $G^h = (V^h, E^h)$ consists of nodes $V^h = V$ from the AHIN G and a **hyperedge** set E^h . Each hyperedge $E_i^h \in E^h$ connects nodes in V^q of an event Q. The center node v_q uniquely identifies a hyperedge.

Thus, an event can be modeled as a hyperedge in the hypergraph, which is capable of modeling the highorder semantics of an event (e.g., multiple authors are co-authored in one paper). An event is abnormal if the event exhibits rare interaction patterns. As depicted in Figure 1(d), the publication events are modeled by hyperedges, which involve multiple types of nodes (i.e., paper, author, and venue), and the abnormal event contains the semantics that data mining specialists cooperate with a radiologist on a COVID-19 paper, which is rarely happened and considered as an anomaly.

Definition 4. Abnormal Event Detection in AHIN. Given an AHIN G = (V, E, X) and an event set $\mathcal{Q} = \{Q_1, \ldots, Q_m\}$, the goal is to learn an anomaly score function $f_{\theta}(\cdot)$ to obtain the anomaly score $s_i = f_{\theta}(Q_i)$ for $\forall Q_i \in \mathcal{Q}$. By choosing the optimal threshold t of anomaly score, the abnormal event set $\mathcal{Q}' = \{Q_i | Q_i \in \mathcal{Q}, f_{\theta}(Q_i) \geq t\}$ can be detected.

4 The Proposed Model

In this section, we present AEHCL, a novel hypergraph contrastive learning method for abnormal event detection. The overall architecture is depicted in Figure 2. To fully capture abnormal event patterns in AHIN, we design contrastive modules at the intra-event and interevent levels. For intra-event, we model the node-pair similarity within an event to capture abnormal interactions of two nodes. High-order abnormal interactions are further captured by the multivariate contrastive task between the center and context nodes. For inter-event, we define the neighbors of events and design contrastive tasks among neighbor events. Based on these contrastive tasks, a novel abnormal event score function is proposed to measure the overall abnormal degrees.

4.1 Type-specific Node Transformation Since an event contains multiple types of nodes, directly using the original features for downstream tasks will suffer from reduced performance [13]. A solution is performing heterogeneous graph convolution [13]. However, the aggregation operation may harm the original feature interaction patterns. Therefore, for a node v_i with type ϕ_i , we directly transform its original feature x_i into a shared latent space via a type-specific transformation matrix W_{ϕ_i} , i.e., $z_i = \sigma(W_{\phi_i} \cdot x_i + b_{\phi_i})$, where $b_{\phi_i} \in \mathbb{R}^d$ is a vector bias. $\sigma(\cdot)$ denotes activation function. After transformation, the representation of each node lies in a shared space with d dimension.



Figure 2: The overall architecture of AEHCL.

4.2 Intra-event Contrastive Module The intraevent contrastive module aims to capture abnormal patterns within an event. These patterns are presented by complicated interactions between involved nodes, which can be classified into two types: pair-wise and multivariate interaction patterns. Accordingly, we design pairwise and multivariate contrastive modules to capture these patterns.

4.2.1 Pair-wise Contrastive Module Pair-wise proximity has been used in hypergraph representation learning methods [19]. The basic intuition behind them is that the pair-wise node matching degree within a hyperedge (event) should be higher than others, which inspires us to model the normal matching patterns of node pairs. Then the node pairs that do not conform to this pattern are considered anomalies. Existing methods directly fuse all pair-wise matching degrees within an event to obtain an event abnormal score. However, this operation may weaken the true abnormal degrees. Therefore, we separately model each node's matching patterns with others rather than fusing them. Specifically, for a node v_j in an event Q_i , we optimize the following contrastive loss:

(4.1)
$$\mathcal{L}_{j}^{pa} = -\log \frac{\sum_{v_{k} \in \mathcal{P}_{j}^{pa}} \exp(sim(z_{j}, z_{k})/\tau)}{\sum_{v_{l} \in \{\mathcal{P}_{j}^{pa} \cup \mathcal{N}_{j}^{pa}\}} \exp(sim(z_{j}, z_{l})/\tau)},$$

where $sim(\cdot)$ is a matching function which is the cosine similarity in our model and τ is the temperature parameter. $\mathcal{P}_{j}^{pa} = \{v_{k} | v_{k} \in V_{i}^{q} \setminus v_{j}\}$ is the positive node set of node v_{j} and we randomly sampled *n* nodes do not belong to Q_{i} as the negative node set \mathcal{N}_{j}^{pa} . Finally, we sum all node's contrastive losses within an event and average losses among all events to get the final pair-wise contrastive loss:

(4.2)
$$\mathcal{L}^{pa} = \frac{1}{m} \sum_{i=1}^{m} \sum_{v_j \in V_i^a} \mathcal{L}_j^{pa}.$$

By minimizing \mathcal{L}^{pa} , the representations of node pairs within an event are similar, i.e., node pairs within a normal event has higher matching degrees compared with abnormal events.

4.2.2Multivariate Contrastive Module The pair-wise contrastive module only captures simple abnormal pair-wise interactions, such as the rare cooperation of two authors. However, there may also exist many complicated multivariate abnormal patterns in the reality. In these cases, the pair-wise interactions within an event are all normal but the event is abnormal when considering interactions between more than two nodes. We model this kind of abnormal event by a multivariate contrastive module. The module captures the multivariate interaction patterns within an event by modeling the compatibility between the center node and context nodes. This compatibility of normal events is higher than abnormal events. For example, a paper's content is highly correlated with the type of published venue and the authors' interests. Specifically, for each node v_i with type ϕ_i , we first add a type embedding o_{ϕ_i} to obtain the type-aware representation h_i , i.e., $h_i = z_i + o_{\phi_i}$. Type embedding makes the model aware of the interactions with heterogeneity to capture more meaningful patterns.

To model the multivariate interactions among context node set V_i^c of event Q_i and obtain the final context representation c_i^{mu} , we apply self-attention [20] followed by a max-pooling operation:

(4.3)
$$c_i^{mu} = \max\text{-pooling}(\operatorname{Att}(\{h_j | v_j \in V_i^c\})),$$

where the max-pooling operation is performed over each dimension of all context nodes' representations. Therefore, c_i^{mu} preserves the most important interaction patterns of context nodes for identifying abnormal events.

After that, we use a simple bilinear scoring function to model the compatibility between the center node and context nodes:

(4.4)
$$s_i^{mu} = \sigma(h_i^{q^{\top}} W_{mu} c_i^{mu}),$$

where h_i^q is the center node representation of the event Q_i and $\sigma(\cdot)$ is the sigmoid activation function. Here, the score s_i^{mu} of normal events should be close to 1, while abnormal events close to 0. Since there is no prior knowledge about abnormal events, we construct the negative context to mimic incompatibility. Instead of randomly replacing context nodes, we first use Kmeans algorithm to perform node clusters based on original features, then for each type of context nodes, we randomly select one node and replace it with another node with the same type but different clusters. In this way, the constructed negative context is different but not easily distinguishable from the positive context, which enhances the capacity for abnormal detection. Then the negative context representation \tilde{c}_i^{mu} and corresponding bilinear score \tilde{s}_i^m can be obtained similarly using Equation (4.3) and Equation (4.4). Finally, we adopt a standard binary cross-entropy (BCE) loss function:

(4.5)
$$\mathcal{L}^{mu} = -\frac{1}{m} \sum_{i=1}^{m} \log(s_i^{mu}) + \log(1 - \tilde{s}_i^{mu}).$$

Inter-event Contrastive Module Abnormal 4.3event patterns may not only limit to abnormal node interactions within an event but also appear among events. Like incompatibility anomaly between local neighbor nodes [15], we believe that there also exists incompatibility anomaly between local neighbor events. Intuitively, the normal events are prone to have similar semantic meanings with their neighbor events, while the abnormal events are not. Based on this assumption, we design contrastive tasks to model the compatibility between neighbor events. Specifically, we first use an attention layer to obtain the event representation. Given an event Q_i , a type-specific attention parameter $P_{\phi_j} \in \mathbb{R}^{d \times d}$ is applied to each context node $v_j \in V_i^c$ with type ϕ_j to obtain attention keys:

$$(4.6) k_j = P_{\phi_j} \cdot z_j,$$

then the attention weights of context nodes are calculated through the softmax function:

(4.7)
$$\alpha_j = \frac{\exp(\sigma(k_j^\top \cdot z_i^q))}{\sum_{l=1}^{|V_i^c|} \exp(\sigma(k_l^\top \cdot z_i^q))}$$

where z_i^q is the representation of the center node. The context representation can be obtained by a weighted sum of all context node embeddings with the learned attention weights:

(4.8)
$$c_i^{in} = \sum_{j=1}^{|V_i^c|} \alpha_j z_j,$$

then we concatenate the context representation c_i^{in} and center node representation z_i^q to obtain the event representation e_i :

(4.9)
$$e_i = c_i^{in} || z_i^q.$$

Next, we define the neighbor event set \mathcal{P}_i^{in} , i.e., positive sample set of event Q_i . Motivated by [24], if two events are connected by more meta-paths [18], they may have more similar semantic meanings. Formally, the \mathcal{P}_i^{in} is defined as follows:

(4.10)
$$\mathcal{P}_i^{in} = \{Q_j | N_{i,j}^{mp} > t^{pos}\},\$$

where $N_{i,j}^{mp}$ denotes the number of meta-paths between event Q_i and Q_j . That is, when the number of metapaths between two events exceeds threshold t^{pos} , then these two events are positive samples of each other. Similarly, we define the negative sample set \mathcal{N}_i^{in} when the number of shared nodes is less than threshold t^{neg} . With the positive set \mathcal{P}_i^{in} and negative set \mathcal{N}_i^{in} , we define the bilinear scoring function to model the compatibility between neighbor events:

(4.11)
$$s_i^{in} = \sigma(e_i^\top W_{in} e_i^p),$$

where e_i^p is the representation of Q_i 's positive sample which is randomly sampled from \mathcal{P}_i^{in} . The score \tilde{s}_i^{in} of the negative event pair can be obtained in the same way. Then, we have the following inter-event contrastive loss:

(4.12)
$$\mathcal{L}^{in} = -\frac{1}{m} \sum_{i=1}^{m} \log(s_i^{in}) + \log(1 - \tilde{s}_i^{in}).$$

4.4 Optimization and Prediction In the training stage, we optimize the above-mentioned three modules jointly. The overall optimization function is:

(4.13)
$$\mathcal{L} = \alpha * \mathcal{L}^{pa} + \beta * \mathcal{L}^{mu} + \gamma * \mathcal{L}^{in},$$

where α , β , and γ are the trade-off parameters to balance three modules. We will discuss them in detail in experiments. The training flow of AEHCL can be found in the supplementary material. After the model is well trained, we use the following abnormal score function to measure the abnormal degree of an event Q_i :

(4.14)

$$f_{\theta}(Q_i) = -(\alpha * \min_{v_j, v_k \in V_i^q} (sim(z_j, z_k)) + \beta * s_i^{mu} + \gamma * s_i^{in})$$

where α , β , and γ are the same as used in Eq. (4.13). $\min(\cdot)$ operation means that we use the minimal nodepair similarity within an event to measure the pair-wise interaction abnormal degree, since the event is likely abnormal only if one node-pair similarity is relatively small. Another choice is the sum of all node-pair similarity scores, which is infeasible since events may have different numbers of entities. We also test other operations (e.g., average) and detailed results are shown in experiments. s_i^{mu} and s_i^{in} mean that we use bilinear scores of positive pairs to measure the multivariate and inter-event abnormal degrees. The bilinear scores of abnormal events are relatively small since they have poor local compatibility. Finally, we add these three scores to obtain the final abnormal event score $s_i =$ $f_{\theta}(Q_i)$. Since we take negative signs, a large abnormal score s_i indicates the event is more likely abnormal.

5 Experiments¹

5.1 Experimental Setup

Datasets. To evaluate the performance of the proposed method, we employ three real-world AHIN datasets. The academic network dataset Aminer, the movie network dataset IMDB, and the takeaway orders dataset Meituan. The basic information of them is summarized in Table 1.

| Dataset | Center nodes | Context Nodes | # Avg. nodes | # Events | # Abnormal events | |
|---------|-----------------|-------------------------------|-----------------|----------|----------------------|--|
| Aminer | Paper: 20567 | Author: 13541 Venue: 115 | 4.3 | 20567 | 1028 | |
| IMDB | Movie: 4661 | Director: 2240 Actor: 5749 | 5 | 4661 | 233 | |
| Meituan | Order: 24330 | Shop: 2265 User: 6972 | 3 | 24330 | 1216 | |

Table 1: Dataset statistics.

Abnormal Event Injection. In previous works [4, 7], the artificial abnormal events are generated as follows: for each event Q_i , n target nodes are sampled from its context node set, where n is sampled from $\{1, 2, 3\}$. These target nodes are then replaced by other sampled nodes with the same type. However, all baselines suffer catastrophic performance under this injection strategy. It may be that this simple injection method makes abnormal events less discriminative with normal ones due to the complex interactions of attributed nodes

within an event. Therefore, similar to [15], for each target node, we sample other k candidate nodes of the same type and calculate the Euclidean distance between each candidate node's attributes and the target node's. Then, we select the node with the largest Euclidean distance to replace the target node, which makes the injected abnormal events more discriminative.

Baselines. We compare our proposed model with three kinds of baselines: abnormal event detection models (APE [4], AEHE [7]); graph abnormal detection models (CoLA [15], ANEMONE [14]; graph representation learning models (metapath2vec [5], HeCo [24], HeteHG-VAE [6]). We also design variants of AEHCL, including intra-event contrastive module AEHCL_p, multivariate contrastive module AEHCL_m, inter-event contrastive module AEHCL_{intra}, pair-wise contrastive module AEHCL_m, inter-event contrastive module AEHCL_{inter} and combinations of different modules: AEHCL_{p+inter} and AEHCL_{m+inter}.

Implementation Details. For all of the baselines, we randomly run 10 times to report the average results. We tune all hyper-parameters to report the best performance. For abnormal detection models, we test all feasible abnormal event score functions and report the best. For AEHCL, we use Adam optimizer and the learning rate is set as 0.001. For optimizing the pair-wise contrastive module solely, we decay the learning rate for every two epochs. The hidden dimension is set as 64, and the original feature dimensions are 108, 128, and 300 for Aminer, IMDB, and Meituan datasets respectively. The temperature parameter τ is set to 1 for all datasets. The hyperparameters α , β , and γ are set as $\{1, 0.8, 0.2\}, \{1, 0.1, 0.1\}, \text{ and } \{0.5, 1, 0.3\}$ in Aminer, IMDB, and Meituan datasets respectively.

5.2Overall **Performance** Following previous works [7], we use two popular evaluation metrics AP and AUC. The overall results are reported in Table 2. Note that HeCo suffers memory issues in the Meituan dataset due to its full-batch training. We have the following observations: (1) The proposed AEHCL model and its variants outperform all the baselines on all the datasets. The improvements are up to 12.0%in AP and 4.6% in AUC, which demonstrates the effectiveness of our model for abnormal event detection in AHIN by fully considering different abnormal event patterns. (2) Generally, only keeping individual module of AEHCL (e.g., AEHCL_p and AEHCL_m) yields poor results, especially in AP scores, while the combination of different modules can obtain more gains. Although adding the inter-event contrastive module will slightly harm the AUC in IMDB and Meituan dataset, it improves much AP score. It's a trade-off between AP and AUC scores and the weights between modules can

¹More details can be found in the supplementary material

| Datasets | Aminer | | IM | DB | Meituan | | | |
|-------------------------------|-----------------------|-----------------------|-----------------------|--------------------------|-----------------------|-----------------------|--|--|
| Metrics | AP | AUC | AP | AUC | AP | AUC | | |
| Metapath2vec | $19.3 \pm 1.2 - 60.9$ | 73.1±1.3 -13.7 | $16.2 \pm 0.9 - 74.3$ | $57.9 \pm 1.9 - 37.2$ | $39.1 \pm 0.0 - 37.0$ | $85.5 \pm 1.0 - 4.3$ | | |
| HeCo | $21.9 \pm 0.0 -55.6$ | $79.8 \pm 0.0 - 5.8$ | $19.1 \pm 0.2 - 69.7$ | $62.6 \pm 0.7 - 32.1$ | _ | _ | | |
| HeteHG-VAE | 23.0 ± 0.0 -53.3 | $55.2 \pm 0.0 - 34.8$ | $11.2 \pm 0.0 - 69.7$ | $52.5 \pm 0.0 - 35.5$ | $12.8 \pm 1.0 - 79.4$ | $63.4 \pm 1.3 - 29.0$ | | |
| ANEMONE | $37.7 \pm 0.4 - 23.5$ | 82.0±0.3 -3.2 | $55.7 \pm 0.1 - 11.7$ | 92.2 ± 0.0 – | $55.5 \pm 0.1 - 10.6$ | <u>89.3</u> ±0.0 – | | |
| CoLA | $33.2 \pm 0.3 - 32.7$ | $83.5 \pm 0.2 - 1.4$ | 63.1 ± 0.0 – | $91.8 \pm 0.1 - 0.4$ | 62.1 ± 0.0 – | $86.6 \pm 0.1 - 3.0$ | | |
| APE | 49.3 ± 0.0 – | 84.7 ± 0.0 – | $30.0 \pm 0.0 - 52.5$ | $81.2 \pm 0.0 - 12.0$ | $30.7 \pm 0.8 - 50.6$ | $83.0 \pm 1.1 - 7.1$ | | |
| AEHE | $48.1 \pm 0.0 - 2.4$ | 84.6±0.0 -0.1 | $28.3 \pm 0.0 -55.2$ | $76.9 \pm 0.0 - 16.6$ | $40.2 \pm 0.2 - 35.3$ | $85.6 \pm 1.1 - 4.1$ | | |
| AEHCL _{intra} | $52.2 \pm 0.2 + 5.9$ | 88.0±0.1 +3.9 | $65.3 \pm 0.0 + 3.5$ | $95.7{\pm}0.0$ +3.8 | $65.9 \pm 0.2 + 6.1$ | $88.0 \pm 0.1 - 1.5$ | | |
| $AEHCL_p$ | 32.0 ± 0.2 -35.1 | $78.3 \pm 0.1 - 7.6$ | $50.8 \pm 0.1 - 19.5$ | $93.8 {\pm} 0.1 {+} 1.7$ | $33.2 \pm 0.4 - 46.5$ | $91.6{\pm}0.1$ +2.6 | | |
| $AEHCL_m$ | $49.5 \pm 0.0 + 0.4$ | $86.0 \pm 0.0 + 1.5$ | $38.9 \pm 0.1 - 38.4$ | $80.2 \pm 0.0 - 13.0$ | $59.3 \pm 0.1 - 4.5$ | $84.3 \pm 0.0 - 5.6$ | | |
| $AEHCL_{p+inter}$ | $33.2 \pm 0.1 - 32.7$ | $79.4 \pm 0.0 - 6.3$ | 54.0 ± 0.0 -14.4 | $93.6 \pm 0.0 + 1.5$ | $33.0 \pm 0.3 - 46.9$ | 79.8±0.1 -10.6 | | |
| $AEHCL_{m+inter}$ | $50.2 \pm 0.0 + 1.8$ | $86.8 \pm 0.0 + 2.5$ | $42.1 \pm 0.0 - 33.3$ | $85.5 \pm 0.0 - 7.3$ | $61.6 \pm 0.1 - 0.8$ | $89.0 \pm 0.0 - 0.3$ | | |
| AEHCL _{inter} | $22.7 \pm 0.4 - 54.0$ | $77.9 \pm 0.1 - 5.7$ | $12.7 \pm 0.2 - 79.9$ | $73.5 \pm 0.0 - 20.3$ | $34.6 \pm 0.2 - 44.3$ | 76.5±0.2 -14.3 | | |
| AEHCL | 55.2±0.1 +12.0 | 88.6±0.1 +4.6 | $67.2{\pm}0.0$ +6.5 | $95.1 \pm 0.0 + 3.1$ | 66.3±0.1 +6.8 | $90.8 \pm 0.0 + 1.7$ | | |

Table 2: Overall performance $(\% \pm \sigma)$ of AEHCL on abnormal event detection on three datasets. The best result of all models is in bold and the best result of previous models is underlined. The improvement (%) is based on the underlined results.

Table 3: Abnormal event cases detected by AEHCL on Aminer and Meituan datasets (dm: data mining).

| Dataset | Abnormal events |
|---------|---|
| Aminer | Paper: Text processing (COVID-19) Author(s): Jiawei Han (dm), Xuan Wang (dm), Kang Zhou (dm) Conf: BIBM (bio-medicine) |
| Meituan | Order A: price > 50, taste id=1 Shop B: category id=1 User C: price (avg.) < 29, taste id (avg.)=2, category id=3 |

be flexibly adjusted under different conditions. In brief, different modules are complementary to each other and focus on capturing different patterns of abnormal events. (3) Graph representation learning methods perform poorly on all datasets, indicating that pure graph representation learning is far from enough to capture abnormal event patterns. The contrastive learning method, HeCo, also performs worse, which shows the challenge of applying contrastive learning to abnormal event detection. The performance of existing abnormal node/event detection models significantly declines in one or two datasets, since they only focus on capturing simple patterns of anomalies (e.g., abnormal attributes of the individual node or abnormal pair-wise interactions).

5.3 Case Study We conduct a case study to show the detection results of AEHCL in Table 3. Note that we haven't reported the abnormal events in IMDB since we only have node features but no actual node information.

In the case of the Aminer dataset, the three authors are all specialized in data mining. However, they collaborate to publish a text-processing paper about COVID-19 at a bio-medicine venue. We found that their regular pattern of this kind of collaboration is at least one biologist co-author. Therefore, though the pair-wise interactions in this event are normal, the multivariate interaction may be abnormal. In the case of Meituan dataset, the price of order A is greater than 50 and the food taste has a tag 1. However, the order is made by a user C whose historical average order price is less than 29 and the average taste tag is 2, which is an uncommon event. Also, user C usually orders food with category 3, but the shop in this order has category 1. In short, this is a common abnormal pair-wise interaction event.

5.4 Model Analysis In this section, we investigate the effect of different abnormal event score functions, as well as the performance of AEHCL under different hyper-parameters.

Effect of Different Abnormal Score Functions. The design of abnormal event score function can be divided into two parts. One is for the pair-wise contrastive module, and another is for the multivariate and inter-event contrastive module since these two parts use different forms of contrastive loss. For the pair-wise contrastive module, besides the *min* operation of our model, there are three other choices: opposite of average of node pair similarities (avg), the standard deviation of node pair similarities (std), and maximum of each node losses within the event using Equation (4.1) (loss). We keep the other two modules unchanged and report



Figure 3: Parameters sensitivity analysis. (Left three figures: effects of different module weights α , β and γ ; right figure: effects of different number of negative samples in pair-wise contrastive module.)

Table 4: Performance (%) on different abnormal score functions in pair-wise contrastive module.

| Datasets | Aminer | | IMDB | | Meituan | |
|-----------------------|--------|------|------|------|---------|------|
| Metrics | AP | AUC | AP | AUC | AP | AUC |
| AEHCLavg | 49.3 | 87.6 | 66.5 | 94.0 | 67.9 | 91.9 |
| AEHCL _{std} | 52.1 | 87.4 | 39.9 | 86.3 | 67.5 | 91.3 |
| AEHCL _{loss} | 50.6 | 85.3 | 59.9 | 93.2 | 67.0 | 90.5 |
| AEHCL _{min} | 55.2 | 88.6 | 67.2 | 95.1 | 66.3 | 90.8 |

Table 5: Performance (%) on different abnormal score functions in multivariate and inter-event contrastive module.

| Datasets | Aminer | | IMDB | | Meituan | |
|----------------------|--------|------|------|------|---------|------|
| Metrics | AP | AUC | AP | AUC | AP | AUC |
| AEHCL ₊ | 55.2 | 88.6 | 67.2 | 95.1 | 66.3 | 90.8 |
| AEHCL_ | 21.8 | 76.0 | 42.2 | 93.3 | 10.9 | 68.5 |
| AEHCL _{+/-} | 41.3 | 85.0 | 65.7 | 94.8 | 60.0 | 89.8 |

the best results in Table 4. The results show that the $AEHCL_{min}$ is not the best for all datasets (it performs a little worse on the Meituan dataset) but the most robust choice. Compared with others, the $AEHCL_{std}$ performs unstable on different datasets, which demonstrates this score setting is less discriminative in some scenarios.

As for the multivariate and inter-event contrastive module, similar to [15], we compare three abnormal event score methods: AEHCL₊ means that only the scores of positive samples are used to calculate abnormal event score, and AEHCL₋ denotes that only using negative samples. $AEHCL_{+/-}$ combines both positive and negative sample scores. We keep the pair-wise contrastive score unchanged and report the best performance in Table 5. Compared with $AEHCL_{+}$, the performance of AEHCL_ decreases dramatically, indicates the negative sample scores have little difference between normal and abnormal events. This makes sense since the negative samples are all sampled in the same way. Combining two scores also harms the overall performance. We suppose that adding the negative contrastive score may disturb the true degree of event abnormality.

Parameters Sensitivity. We investigate the effects of the three key hyper-parameters (α , β and γ) in AE-HCL. For each parameter, we change it from 0 to 1 with step 0.1 while fixing the other two parameters as the best ones. AP score is reported in Figure 3 (AUC score has a similar trend). There have following observations: (1) The pair-wise contrastive module is significant for abnormal event detection. The model performance consistently increases along with the increment of α on two datasets, and also achieves best when α is near 0.5 on Meituan dataset. This indicates that the pair-wise interaction anomalies may be more general and can be well captured by the pair-wise contrastive module. (2) The multivariate contrastive module has different effects on different datasets. Though an increasing trend of effects can be found when increasing β on Meituan dataset, the effects suffer decreasing on the other two datasets when β is too large, notably on IMDB dataset. The reason may be that these two datasets has fewer multivariate interaction anomalies. (3) The inter-event contrastive module plays an complementary role in detecting abnormal events. The performance has a decreasing trend on all datasets when increasing γ . However, as discussed in Section 5.2, the model performs better when γ is relatively small (e.g., $\gamma = 0.1$) than setting $\gamma = 0$. Therefore, despite the relatively few inter-event anomalies, this module is indispensable for further improving the detection performance.

We also analyze the significance of negative sample numbers n in the pair-wise contrastive module. Figure 3 shows AP score under different number of negative samples. The performance on IMDB and Meituan datasets are relatively stable with different n. however, the model suffers performance decay when n is large (n > 80) on the Aminer dataset. Since we have few limitations on the sampling strategy of negative samples, there may involve some entities whose features are similar to the target entities when n is too large. Therefore, we set n = 10 for all datasets. More sampling strategies are expected to study in further work.

6 Conclusion

In this paper, we explore the challenge problem of abnormal event detection in AHIN. Different from previous works focusing on detecting abnormal events with simple pair-wise patterns, we model events as hyperedges and proposed a hypergraph contrastive learning method, named AEHCL, which includes three contrastive modules to capture pair-wise, multivariate, and inter-event abnormal patterns. Moreover, an abnormal event score function is proposed to measure the abnormal degrees. Extensive experiments on three real-world datasets demonstrate the effectiveness of AEHCL.

7 Acknowledgement

This work is supported in part by the National Natural Science Foundation of China (No. U20B2045, 62192784, 62172052, 62002029, 62172052, U1936014) and BUPT Excellent Ph.D. Students Foundation (No. CX2021118).

References

- L. AKOGLU, H. TONG, J. VREEKEN, AND C. FALOUT-SOS, Fast and reliable anomaly detection in categorical data, in CIKM, ACM, 2012.
- [2] C. BERGE, Hypergraphs combinatorics of finite sets, vol. 45 of North-Holland mathematical library, 1989.
- [3] T. CHEN, S. KORNBLITH, M. NOROUZI, AND G. E. HINTON, A simple framework for contrastive learning of visual representations, in ICML, vol. 119 of Proceedings of Machine Learning Research, PMLR, 2020.
- [4] T. CHEN, L. TANG, Y. SUN, Z. CHEN, AND K. ZHANG, Entity embedding-based anomaly detection for heterogeneous categorical events, in IJCAI, 2016.
- [5] Y. DONG, N. V. CHAWLA, AND A. SWAMI, metapath2vec: Scalable representation learning for heterogeneous networks, in KDD, ACM, 2017.
- [6] H. FAN, F. ZHANG, Y. WEI, Z. LI, C. ZOU, Y. GAO, AND Q. DAI, *Heterogeneous hypergraph variational autoencoder for link prediction*, IEEE Trans. Pattern Anal. Mach. Intell., 44 (2022).
- [7] S. FAN, C. SHI, AND X. WANG, Abnormal event detection via heterogeneous information network embedding, in CIKM, ACM, 2018.
- [8] H. GUI, J. LIU, F. TAO, M. JIANG, B. NORICK, L. M. KAPLAN, AND J. HAN, *Embedding learning with events* in heterogeneous information networks, IEEE Trans. Knowl. Data Eng., 29 (2017).
- [9] X. HAN, H. CHENG, D. XU, AND S. YUAN, Interpretablesad: Interpretable anomaly detection in sequential log data, in IEEE BigData, IEEE, 2021.
- [10] K. HASSANI AND A. H. K. AHMADI, Contrastive multi-view representation learning on graphs, in ICML, PMLR, 2020.

- [11] K. HE, H. FAN, Y. WU, S. XIE, AND R. B. GIRSHICK, Momentum contrast for unsupervised visual representation learning, in CVPR, 2020.
- [12] B. HU, Z. ZHANG, C. SHI, J. ZHOU, X. LI, AND Y. QI, Cash-out user detection based on attributed heterogeneous information network with a hierarchical attention mechanism, in AAAI, AAAI Press, 2019.
- [13] L. HU, T. YANG, C. SHI, H. JI, AND X. LI, Heterogeneous graph attention networks for semi-supervised short text classification, in EMNLP/IJCNLP (1), Association for Computational Linguistics, 2019.
- [14] M. JIN, Y. LIU, Y. ZHENG, L. CHI, Y. LI, AND S. PAN, ANEMONE: graph anomaly detection with multi-scale contrastive learning, in CIKM, ACM, 2021.
- [15] Y. LIU, Z. LI, S. PAN, C. GONG, C. ZHOU, AND G. KARYPIS, Anomaly detection on attributed networks via contrastive self-supervised learning, IEEE Trans. Neural Networks Learn. Syst., 33 (2022).
- [16] Y. LU, Y. FANG, AND C. SHI, Meta-learning on heterogeneous information networks for cold-start recommendation, in KDD, ACM, 2020.
- [17] C. PARK, D. KIM, J. HAN, AND H. YU, Unsupervised attributed multiplex network embedding, in AAAI, AAAI Press, 2020.
- [18] C. SHI, Y. LI, J. ZHANG, Y. SUN, AND P. S. YU, A survey of heterogeneous information network analysis, IEEE Trans. Knowl. Data Eng., 29 (2017).
- [19] K. TU, P. CUI, X. WANG, F. WANG, AND W. ZHU, Structural deep embedding for hyper-networks, in AAAI, AAAI Press, 2018.
- [20] A. VASWANI, N. SHAZEER, N. PARMAR, J. USZKO-REIT, L. JONES, A. N. GOMEZ, L. KAISER, AND I. POLOSUKHIN, Attention is all you need, in NIPS, 2017.
- [21] P. VELICKOVIC, W. FEDUS, W. L. HAMILTON, P. LIÒ, Y. BENGIO, AND R. D. HJELM, *Deep graph infomax*, in ICLR (Poster), OpenReview.net, 2019.
- [22] H. WANG, C. ZHOU, J. WU, W. DANG, X. ZHU, AND J. WANG, *Deep structure learning for fraud detection*, in ICDM, IEEE Computer Society, 2018.
- [23] X. WANG, H. JI, C. SHI, B. WANG, Y. YE, P. CUI, AND P. S. YU, *Heterogeneous graph attention network*, in WWW, ACM, 2019.
- [24] X. WANG, N. LIU, H. HAN, AND C. SHI, Selfsupervised heterogeneous graph neural network with cocontrastive learning, in KDD, ACM, 2021.
- [25] Z. WANG, Z. CHEN, J. NI, H. LIU, H. CHEN, AND J. TANG, Multi-scale one-class recurrent neural networks for discrete event sequence anomaly detection, in KDD, ACM, 2021.
- [26] M. ZHENG, C. ZHOU, J. WU, S. PAN, J. SHI, AND L. GUO, Fraudne: a joint embedding approach for fraud detection, in IJCNN, IEEE, 2018.
- [27] Y. ZHENG, M. JIN, Y. LIU, L. CHI, K. T. PHAN, AND Y.-P. P. CHEN, Generative and contrastive selfsupervised learning for graph anomaly detection, IEEE Trans. Knowl. Data Eng., (2021).