



Dynamic Graph Information Bottleneck

Haonan Yuan

School of Computer Science and
Engineering, BDBC,
Beihang University
Beijing, China
yuanhn@buaa.edu.cn

Qingyun Sun*

School of Computer Science and
Engineering, BDBC,
Beihang University
Beijing, China
sunqy@buaa.edu.cn

Xingcheng Fu

Key Lab of Education Blockchain and
Intelligent Technology,
Guangxi Normal University
Guilin, Guangxi, China
fuxc@gxnu.edu.cn

Cheng Ji

School of Computer Science and
Engineering, BDBC,
Beihang University
Beijing, China
jicheng@buaa.edu.cn

Jianxin Li

School of Computer Science and
Engineering, BDBC,
Beihang University
Beijing, China
lijx@buaa.edu.cn

ABSTRACT

Dynamic Graphs widely exist in the real world, which carry complicated spatial and temporal feature patterns, challenging their representation learning. Dynamic Graph Neural Networks (DGNNs) have shown impressive predictive abilities by exploiting the intrinsic dynamics. However, DGNNs exhibit limited robustness, prone to adversarial attacks. This paper presents the novel *Dynamic Graph Information Bottleneck (DGIB)* framework to learn robust and discriminative representations. Leveraged by the Information Bottleneck (IB) principle, we first propose the expected optimal representations should satisfy the *Minimal-Sufficient-Consensual (MSC)* Condition. To compress redundant as well as conserve meritorious information into latent representation, DGIB iteratively directs and refines the structural and feature information flow passing through graph snapshots. To meet the *MSC* Condition, we decompose the overall IB objectives into $DGIB_{MS}$ and $DGIB_C$, in which the $DGIB_{MS}$ channel aims to learn the minimal and sufficient representations, with the $DGIB_C$ channel guarantees the predictive consensus. Extensive experiments on real-world and synthetic dynamic graph datasets demonstrate the superior robustness of DGIB against adversarial attacks compared with state-of-the-art baselines in the link prediction task. To the best of our knowledge, DGIB is the first work to learn robust representations of dynamic graphs grounded in the information-theoretic IB principle.

CCS CONCEPTS

• **Mathematics of computing** → **Graph algorithms**; • **Computing methodologies** → **Neural networks**; **Learning latent representations**.

*Corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
WWW '24, May 13–17, 2024, Singapore, Singapore
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0171-9/24/05
<https://doi.org/10.1145/3589334.3645411>

KEYWORDS

dynamic graph neural networks, robust representation learning, information bottleneck

ACM Reference Format:

Haonan Yuan, Qingyun Sun, Xingcheng Fu, Cheng Ji, and Jianxin Li. 2024. Dynamic Graph Information Bottleneck. In *Proceedings of the ACM Web Conference 2024 (WWW '24), May 13–17, 2024, Singapore, Singapore*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3589334.3645411>

1 INTRODUCTION

Dynamic graphs are prevalent in real-world scenarios, encompassing domains such as social networks [4], financial transaction [54], and traffic networks [55], *etc.* Due to their intricate spatial and temporal correlation patterns, addressing a wide spectrum of applications like web link co-occurrence analysis [6, 17], relation prediction [53], anomaly detection [5] and traffic flow analysis [24], *etc.* poses significant challenges. Leveraging their exceptional expressive capabilities, dynamic graph neural networks (DGNNs) [14] intrinsically excel in the realm of dynamic graph representation learning by modeling both spatial and temporal predictive patterns, which is achieved with the combined merits of graph neural networks (GNNs)-based models and sequence-based models.

Recently, there has been an increasing emphasis on the efficacy enhancement of the DGNNs [10, 11, 39, 56, 57], with a specific focus on augmenting their capabilities to capture intricate spatio-temporal feature patterns against the first-order Weisfeiler-Leman (1-WL) isomorphism test [47]. However, most of the existing works still struggle with the over-smoothing phenomenon brought by the message-passing mechanism in most (D)GNNs [7]. Specifically, the inherent node features contain potentially irrelevant and spurious information, which is aggregated over the edges, compromising the resilience of DGNNs, and prone to ubiquitous noise of in-the-wild testing samples with possible adversarial attacks [38, 45, 49, 50].

To against adversarial attacks, the Information Bottleneck (IB) [42, 43] introduces an information-theoretic theory for robust representation learning, which encourages the model to acquire the most informative and predictive representation of the target, satisfying both the minimal and sufficient (*MS*) assumption. In Figure 1(a), the IB principle encourages the model to capture the maximum




Figure 1: Comparison among different IB principles.

mutual information about the target and make accurate predictions (*Sufficient*). Simultaneously, it discourages the inclusion of redundant information from the input that is unrelated to predicting the target (*Minimal*). By adhering to this paradigm, the trained model naturally mitigates overfitting and becomes more robust to potential noise and adversarial attacks. However, directly applying IB to the scene of dynamic graphs faces significant challenges. First, IB necessitates that the learning process adheres to the Markovian dependence with independent and identically distributed (*i.i.d.*) data, which is inherently not satisfied for the non-Euclidean dynamic graphs. Second, dynamic graphs exhibit coupled structural and temporal features, where these discrete and intricate information flows make optimizing the IB objectives intractable.

A few attempts have been made to extend the IB principle to *static* graphs [37, 48, 51, 52] by a two-step paradigm (Figure 1(b)), which initially obtains tractable representations by GNNs, and subsequently models the distributions of variables by variational inference [43], which enables the estimation of IB objectives. However, as the crucial structural information is absent from the direct IB optimization, it leads to unsatisfying robust performance. To make structures straightforwardly involved in the optimization process, [46] explicitly compresses the input from both the graph structure and node feature perspectives iteratively in a tractable and constrained searching space, which is then optimized with the estimated variational bounds and leads to better robustness (Figure 1(c)). Accordingly, the learned representations satisfy the *MS* Assumption, which can reduce the impact of structural and intricate feature noise, alleviate overfitting, and strengthen informative and discriminative prediction ability. To the best of our knowledge, there’s no successful work to adapt the IB principle to *dynamic* graph representation learning with spatio-temporal graph structures directly involved.

Due to the local (single snapshot) and global (overall snapshots) joint reliance for prediction tasks on dynamic graphs, which introduce complex correlated spatio-temporal patterns evolving with time, extending the intractable IB principle to the dynamic graph representation learning is a non-trivial problem confronting the following major challenges:

- How to understand what constitutes the optimal representation that is both discriminative and robust for downstream task prediction under the dynamic scenario? (> Section 4.1)
- How to appropriately compress the input dynamic graph features by optimizing the information flow across graph snapshots with structures straightforwardly involved? (> Section 4.2)
- How to optimize the intractable IB objectives, which are incalculable on the non-Euclidean dynamic graphs? (> Section 4.3)

Present work. To tackle the aforementioned challenges, we propose the innovative *Dynamic Graph Information Bottleneck (DGIB)* framework for robust dynamic graph representation learning, aiming at striking a balance between the expressiveness and robustness of DGNNs. To understand what contributes to the optimal representations that are both informative and robust for prediction in dynamic scenarios, we propose the *Minimal-Sufficient-Consensual (MSC)* Condition with empirical and theoretical analysis. To conserve the informative features, we design the spatio-temporal sampling mechanism to derive the local dependence assumption, based on which we establish the DGIB principle that guides the information flow to iteratively refine structures and features. Specifically, we decompose the overall DGIB objective into $DGIB_{MS}$ and $DGIB_C$ channels, cooperating to jointly satisfy the proposed *MSC* Condition. To make the IB objectives tractable, we introduce the variational upper and lower bounds for optimizing $DGIB_{MS}$ and $DGIB_C$. We highlight the advantages of DGIB as follows:

- We propose the novel DGIB framework for robust dynamic graph representation learning. To the best of our knowledge, this is the first exploration to extend IB on dynamic graphs with structures directly involved in the IB optimization.
- We investigate a new insight and propose the *Minimal-Sufficient-Consensual (MSC)* Condition, which can be satisfied by the cooperation of both $DGIB_{MS}$ and $DGIB_C$ channels to refine the spatio-temporal information flow for feature compression. We further introduce their variational bounds for optimization.
- Extensive experiments on both real-world and synthetic dynamic graph datasets demonstrate the superior robustness of our DGIB against targeted and non-targeted adversarial attacks compared with state-of-the-art baselines.

2 RELATED WORK

2.1 Dynamic Graph Representation Learning

Dynamic Graphs find applications in a wide variety of disciplines, including social networks, recommender systems, epidemiology, *etc.* Our primary concern lies in the representation learning for the common *discrete* dynamic graphs, which encompass multiple discrete graph snapshots arranged in chronological order.

Dynamic Graph Neural Networks (DGNNs) are widely adopted to learn dynamic graph representations by intrinsically modeling both spatial and temporal predictive patterns, which can be divided into two categories [14]. (1) **Stacked DGNNs** employ separate GNNs to process each graph snapshot and forward the output of each GNN to deep sequential models. Stacked DGNNs alternately model the dynamics to learn representations, which are the mainstream. (2) **Integrated DGNNs** function as encoders that combine GNNs and deep sequential models within a single layer, unifying spatial and temporal modeling. The deep sequential models are applied to initialize the weights of GNN layers.

However, as dynamic graphs are often derived from open environments, which inherently contain noise and redundant features unrelated to the prediction task, the DGNN performance in downstream tasks is compromised. Additionally, DGNNs are susceptible to inherent over-smoothing issues, making them sensitive and vulnerable to adversarial attacks. Currently, no robust DGNN solutions have been effectively proposed.

2.2 Information Bottleneck

The information Bottleneck (IB) principle aims to discover a concise code for the input signal while retaining the maximum information within the code for signal processing [42]. [43] initially extends Variational Information Bottleneck (VIB) to deep learning, facilitating diverse downstream applications in fields such as reinforcement learning [9, 16], computer vision [13, 23], and natural language processing [29], *etc.* IB is employed to select a subset of input features, such as pixels in images or dimensions in embeddings that are maximally discriminative concerning the target. Nevertheless, research on IB in the non-Euclidean dynamic graphs has been relatively limited, primarily due to its intractability of optimizing.

There are some prior works extend IB on *static* graphs, which can be categorized into two groups based on whether the graph structures are straightforwardly involved in the IB optimization process. **(1) Non-structure-involved**, which follows the two-step learning paradigm illustrated in Figure 1(b). SIB [52] is proposed for the critical subgraph recognition problem. HGIB [51] implements the consensus hypothesis of heterogeneous information networks in an unsupervised manner. VIB-GSL [37] first leverages IB to graph structure learning, *etc.* To make IB objectives tractable, they model the distributions of representations by variational inference, which enables the estimation of the IB terms. **(2) Structure-involved** illustrated in Figure 1(c). To make structures directly involved in the feature compression process, the pioneering GIB [46] explicitly extracts information with regularization of the structure and feature information in a tractable searching space.

In conclusion, structure-involved GIBs demonstrate superior advantages over non-structure-involved ones by leveraging the significant graph structure patterns, which can reduce the impact of environment noise, enhance the robustness of models, as well as be informative and discriminative for downstream prediction tasks. However, extending IB to dynamic graph learning proves challenging, as the impact of the spatio-temporal correlations within the Markovian process should be elaborately considered.

3 NOTATIONS AND PRELIMINARIES

In this paper, random variables are denoted as **bold** letters while their realizations are *italic* letters. The ground-truth distribution is represented as $\mathbb{P}(\cdot)$, and $\mathbb{Q}(\cdot)$ denotes its approximation.

Notation. We primarily consider the *discrete* dynamic representation learning. A discrete dynamic graph can be denoted as a series of graphs snapshots $\text{DG} = \{\mathcal{G}^t\}_{t=1}^T$, where T is the time length. $\mathcal{G}^t = (\mathcal{V}^t, \mathcal{E}^t)$ is the graph at time t , where \mathcal{V}^t is the node set and \mathcal{E}^t is the edge set. Let $\mathbf{A}^t \in \{0, 1\}^{N \times N}$ be the adjacency matrix and $\mathbf{X}^t \in \mathbb{R}^{N \times d}$ be the node features, where $N = |\mathcal{V}^t|$ denotes the number of nodes and d denotes the feature dimensionality.

Dynamic Graph Representation Learning. As the most challenging task of dynamic graph representation learning, the future link prediction aims to train a model $f_\theta : \mathcal{V} \times \mathcal{V} \mapsto \{0, 1\}^{N \times N}$ that predicts the existence of edges at $T + 1$ given historical graphs $\mathcal{G}^{1:T}$ and next-step node features \mathbf{X}^{T+1} . Concretely, $f_\theta = w \circ g$ is compound of a DGNN $w(\cdot)$ to learn node representations and a link predictor $g(\cdot)$ for link prediction, *i.e.*, $\mathbf{Z}^{T+1} = w(\mathcal{G}^{1:T}, \mathbf{X}^{T+1})$ and $\hat{\mathbf{Y}}^{T+1} = g(\mathbf{Z}^{T+1})$. Our goal is to learn a robust representation against adversarial attacks with the optimal parameters θ^* .

Information Bottleneck. The Information Bottleneck (IB) principle trades off the data fit and robustness using mutual information (MI) as the cost function and regularizer. Given the input \mathbf{X} , representation \mathbf{Z} of \mathbf{X} and target \mathbf{Y} , the tuple $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ follows the Markov Chain $\langle \mathbf{Y} \rightarrow \mathbf{X} \rightarrow \mathbf{Z} \rangle$. IB learns the minimal and sufficient representation \mathbf{Z} by optimizing the following objective:

$$\mathbf{Z} = \arg \min_{\mathbf{Z}} -I(\mathbf{Y}; \mathbf{Z}) + \beta I(\mathbf{X}; \mathbf{Z}), \quad (1)$$

where β is the Lagrangian parameter to balance the two terms, and $I(\mathbf{X}; \mathbf{Y})$ represents the mutual information between \mathbf{X} and \mathbf{Y} .

4 DYNAMIC GRAPH INFORMATION BOTTLENECK

In this section, we elaborate on the proposed DGIB, where its principle and framework are shown in Figure 2. First, we propose the *Minimal-Sufficient-Consensual (MSC)* Condition that the expected optimal representations should satisfy. Then, we derive the DGIB principle by decomposing it into DGIB_{MS} and DGIB_C channels, which cooperate and contribute to satisfying the proposed *MSC* Condition. Lastly, we instantiate the DGIB principle with tractable variational bounds for efficient IB objective optimization.

4.1 DGIB Optimal Representation Condition

Given the input \mathbf{X} and label \mathbf{Y} , the sufficient statistics theory [36] identifies the optimal representation of \mathbf{X} , namely, $S(\mathbf{X})$, which effectively encapsulates all the pertinent information contained within \mathbf{X} concerning \mathbf{Y} , namely, $I(S(\mathbf{X}); \mathbf{Y}) = I(\mathbf{X}; \mathbf{Y})$. Steps further, the minimal and sufficient statistics $T(\mathbf{X})$ establish a fundamental sufficient partition of \mathbf{X} . This can be expressed through the Markov Chain: $\text{MC}_{\text{IB}} \langle \mathbf{Y} \rightarrow \mathbf{X} \rightarrow S(\mathbf{X}) \rightarrow T(\mathbf{X}) \rangle$, which holds true for a minimal sufficient statistics $T(\mathbf{X})$ in conjunction with any other sufficient statistics $S(\mathbf{X})$. Leveraging the Data Processing Inequality (DPI) [3], the *Minimal* and *Sufficient (MS)* Assumption satisfied representation can be optimized by:

$$T(\mathbf{X}) = \arg \min_{S(\mathbf{X}): I(S(\mathbf{X}); \mathbf{Y}) = I(\mathbf{X}; \mathbf{Y})} I(S(\mathbf{X}); \mathbf{X}), \quad (2)$$

where the mapping $S(\cdot)$ can be relaxed to any encoder $\mathbb{P}(T | \mathbf{X})$, which allows $S(\cdot)$ to capture as much as possible of $I(\mathbf{X}; \mathbf{Y})$.

Case Study. The structure-involved GIB adapts the MC_{IB} to $\text{MC}_{\text{GIB}} \langle \mathbf{Y} \rightarrow \mathcal{G} \rightarrow S(\mathcal{G}) \rightarrow \mathbf{Z} \rangle$, where \mathcal{G} contains both structures (\mathbf{A}) and node features (\mathbf{X}), leading to the minimal and sufficient \mathbf{Z} with a balance between expressiveness and robustness. However, we observe a counterfactual result that if directly adapting MC_{GIB} to dynamic graphs as $\text{MC}_{\text{DGIB}} \langle \mathbf{Y}^{T+1} \rightarrow (\mathcal{G}^{1:T}, \mathbf{X}^{T+1}) \rightarrow S(\mathcal{G}^{1:T}, \mathbf{X}^{T+1}) \rightarrow \mathbf{Z}^{T+1} \rangle$ and optimizing the overall objective:

$$\mathbf{Z}^{T+1} = \arg \min_{\mathbf{Z}^{T+1}} \left[-I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) + \beta I(\mathcal{G}^{1:T}, \mathbf{X}^{T+1}; \mathbf{Z}^{T+1}) \right] \quad (3)$$

will lead to sub-optimal *MS* representation $\hat{\mathbf{Z}}^{T+1}$.

Following settings in Appendix C.2, we adapt the GIB [46] to dynamic scenarios by processing each individual graph with original GIB [46] and then aggregating each output \mathbf{Z}^t with LSTM [15] to learn the comprehensive representation $\hat{\mathbf{Z}}^{T+1}$, where the two steps are *jointly* optimized by Eq. (3). We evaluate the model performance of the link prediction task on the dynamic graph dataset ACT [22]. As depicted in Figure 3, the results reveal a noteworthy finding:




Figure 2: The proposed DGIB principle and its overall framework. (a) simultaneously maximizing the mutual information between the representation and the target while constraining information with the input graphs. The significant graph structures are directly involved in the optimizing process. (b) iteratively compresses structures and node features between graphs. The overall $\mathcal{L}_{\text{DGIB}}$ is decomposed to DGIB_{MS} and DGIB_{C} channels, which act jointly to satisfy the MSC Condition.




Figure 3: Case study between adapted GIB and DGIB-Cat.

the prediction performance achieved by $\hat{\mathbf{Z}}^{t+1}$ for the next time-step graph \mathcal{G}^{t+1} during training significantly *surpasses* that of $\hat{\mathbf{Z}}^{t+1}$ for the target graph \mathcal{G}^{T+1} during validating and testing.

We attribute the above results to the *laziness* of deep neural networks [25, 40], leading to *non-consensual* feature compression for each time step. Specifically, there exist $T+1$ independent and identically distributed Markov Chains $MC_{\text{DGIB}}^t < \mathbf{Y}^{t+1} \rightarrow (\mathcal{G}^t, \mathbf{X}^{t+1}) \rightarrow S(\mathcal{G}^t, \mathbf{X}^{t+1}) \rightarrow \mathbf{Z}^{t+1}$. Optimizing each MC_{DGIB}^t as:

$$\mathbf{Z}^{t+1} = \arg \min_{\mathbf{Z}^{t+1}} [-I(\mathbf{Y}^{t+1}; \mathbf{Z}^{t+1}) + \beta I(\mathcal{G}^t, \mathbf{X}^{t+1}; \mathbf{Z}^{t+1})] \quad (4)$$

guarantees the learned \mathbf{Z}^{t+1} meeting the *MS* Assumption respectively for each \mathcal{G}^{t+1} . However, as the spatio-temporal patterns among graph snapshots are intricately complicated, their dependencies are coupled. The adapted GIB tends to converge to a sub-optimal *trivial* solution that satisfies each MC_{DGIB}^t , such that every $\hat{\mathbf{Z}}^{t+1}$ is close to *MS* status for the next-step label \mathbf{Y}^{t+1} , but fail to predict the final \mathbf{Y}^{T+1} due to its laziness. Consequently, the objective in Eq. (3) degrades to simply optimize the union of $\{\hat{\mathbf{Z}}^t\}_{t=1}^{T+1}$.

The above analysis confirms the fact that the global representation $\hat{\mathbf{Z}}^{T+1}$ does not align closely with the *MS* Assumption. To encourage $\hat{\mathbf{Z}}^{T+1}$ to reach the *global* minimal and sufficient status, we apply an additional *Consensual* Constraint $C(\theta)$ on the sequence of MC_{DGIB}^t , which acts like the “baton” for compression process, and encourages conserving the consensual parts. We conclude the *Minimal-Sufficient-Consensual (MSC)* Condition as follows.

Assumption 1 (Minimal-Sufficient-Consensual (MSC) Condition). Given $\text{DG} = \{\mathcal{G}^t\}_{t=1}^T$, the optimal representation \mathbf{Z}^{T+1} for the robust future link prediction should satisfy the *Minimal-Sufficient-Consensual (MSC)* Condition, such that:

$$\mathbf{Z}^{T+1} = \arg \min_{S(\mathcal{D}): I(S(\mathcal{D}); \mathbf{Y}^{T+1}) = I(\mathcal{D}; \mathbf{Y}^{T+1})} I(S(\mathcal{D}); \mathcal{D}), \quad (5)$$

where \mathcal{D} is the training data containing previous graphs $\mathcal{G}^{1:T}$ and node feature \mathbf{X}^{T+1} at the next time-step, and $S(\cdot)$ is the partitions of \mathcal{D} , which is implemented by a stochastic encoder $\mathbb{P}(\mathbf{Z}^{T+1} | \mathcal{D}, C(\theta))$, where $C(\theta)$ satisfies the *Consensual* Constraint.

Assumption 1 declares the optimal representation \mathbf{Z}^{T+1} for the robust future link prediction task of dynamic graphs should be minimal, sufficient and consensual (*MSC*).

4.2 DGIB Principle Derivation

Definition 1 (Dynamic Graph Information Bottleneck). Given $\text{DG} = \{\mathcal{G}^t\}_{t=1}^T$, and the nodes feature \mathbf{X}^{T+1} at the next time-step, the Dynamic Graph Information Bottleneck (DGIB) is to learn the optimal representation \mathbf{Z}^{T+1} that satisfies *MSC* Condition by:

$$\begin{aligned} \mathbf{Z}^{T+1} &= \arg \min_{\mathbb{P}(\mathbf{Z}^{T+1} | \mathcal{D}, C(\theta)) \in \Omega} \text{DGIB}(\mathcal{D}, \mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) \\ &\triangleq \left[-I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) + \beta I(\mathcal{D}; \mathbf{Z}^{T+1}) \right], \end{aligned} \quad (6)$$

where $\mathcal{D} = (\mathcal{G}^{1:T}, \mathbf{X}^{T+1})$ is the input training data, and Ω defines the search space of the optimal DGNN model $\mathbb{P}(\mathbf{Z}^{T+1} | \mathcal{D}, C(\theta))$.

Compared with Eq. (3), Eq. (6) satisfies the additional *Consensual* Constraint $C(\theta)$ on $\mathbb{P}(\mathbf{Z}^{T+1} | \mathcal{D})$ within the underlying search space Ω . This encourages \mathbf{Z}^{T+1} to align more effectively with the *MSC* Condition. Subsequently, we decompose the overall DGIB into DGIB_{MS} and DGIB_{C} channels, both sharing the same IB structure as Eq. (1). DGIB_{MS} aims to optimize \mathbf{Z}^{T+1} under $\mathbb{P}(\mathbf{Z}^{T+1} | \mathcal{D}, \theta)$, while DGIB_{C} encourages $\mathbf{Z}^{1:T}$ and \mathbf{Z}^{T+1} share consensual predictive pattern for predicting \mathbf{Y}^{T+1} under $\mathbb{P}(\mathbf{Z}^{T+1} | \mathbf{Z}^{1:T}, C(\theta))$. DGIB_{MS} and DGIB_{C} cooperate and contribute to satisfying the proposed *MSC* Condition under $\mathbb{P}(\mathbf{Z}^{T+1} | \mathbf{Z}^{1:T}, C(\theta))$.

4.2.1 Deriving the DGIB_{MS}. Directly optimizing Eq. (6) poses a significant challenge due to the complex spatio-temporal correlations. The *i.i.d.* assumption, typically necessary for estimating variational bounds, is a key factor in rendering the optimization of the IB objectives feasible. However, the *i.i.d.* assumption cannot be reached in dynamic graphs, as node features exhibit correlations owing to the underlying graph structures, and the dependencies are intricate. To solve the above challenges, we propose the Spatio-Temporal Local Dependence Assumption following [34].

Assumption 2 (Spatio-Temporal Local Dependence). Given $\text{DG} = \{\mathcal{G}^t\}_{t=1}^T$, let $\mathcal{N}_{ST}(v, k, t)$ be the spatio-temporal k -hop neighbors of any node $v \in \mathcal{V}$. The rest of the DG will be independent of node v and its spatio-temporal k -hop neighbors, *i.e.*:

$$\mathbb{P}\left(\mathbf{X}_v^t \mid \mathcal{G}_{\mathcal{N}_{ST}(v,k,t)}^{t-1:t}, \overline{\mathcal{G}}_{\{v\}}^{1:T}\right) = \mathbb{P}\left(\mathbf{X}_v^t \mid \mathcal{G}_{\mathcal{N}_{ST}(v,k,t)}^{t-1:t}\right), \quad (7)$$

where $\mathcal{G}_{\mathcal{N}_{ST}(v,k,t)}^{t-1:t}$ denotes $\mathcal{N}_{ST}(v, k, t)$ -related subgraphs, and $\overline{\mathcal{G}}_{\{v\}}^{1:T}$ denotes complement graphs in terms of node v and associated edges.

Assumption 2 is applied to constrain the search space Ω as $\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathcal{D}, \theta)$, which leads to a more feasible DGIB_{MS}:

$$\begin{aligned} \mathbf{Z}^{T+1} &= \arg \min_{\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathcal{D}, \theta) \in \Omega} \text{DGIB}_{MS}(\mathcal{D}, \mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) \\ &\triangleq \left[- \underbrace{I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1})}_{\text{Eq. (15)}} + \beta_1 \underbrace{I(\mathcal{D}; \mathbf{Z}^{T+1})}_{\text{Eq. (17)/(18), Eq. (19)}} \right]. \end{aligned} \quad (8)$$

Conv Layer. We assume the learning process follows the Markovian dependency in Figure 2. During each iteration l , \mathbf{X}_v^l will be updated by incorporating its spatio-temporal neighbors $\mathcal{N}_{ST}(v, k, t)$ on the refined structure $\hat{\mathbf{A}}^t$, which controls the information flow across graph snapshots. In this way, DGIB_{MS} requires to optimize the distributions of $\mathbb{P}(\hat{\mathbf{A}}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \mathbf{A}^t)$ and $\mathbb{P}(\mathbf{Z}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \hat{\mathbf{A}}^t)$.

Variational Bounds of DGIB_{MS}. We respectively introduce the lower bound of $I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1})$ with respect to [31], and the upper bound of $I(\mathcal{D}; \mathbf{Z}^{T+1})$ inspired by [46], for Eq. (8).

Proposition 1 (Lower Bound of $I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1})$).

$$\begin{aligned} I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) &\geq 1 + \mathbb{E}_{\mathbb{P}(\mathbf{Y}^{T+1}, \mathbf{Z}^{T+1})} \left[\log \frac{\mathbb{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathbb{Q}_2(\mathbf{Y}^{T+1})} \right] \\ &\quad - \mathbb{E}_{\mathbb{P}(\mathbf{Y}^{T+1})\mathbb{P}(\mathbf{Z}^{T+1})} \left[\frac{\mathbb{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathbb{Q}_2(\mathbf{Y}^{T+1})} \right]. \end{aligned} \quad (9)$$

Proposition 2 (Upper Bound of $I(\mathcal{D}; \mathbf{Z}^{T+1})$). Let $\mathcal{I}_A, \mathcal{I}_Z \subset [T+1]$ be the random time indices. Based on the Markov property $\mathcal{D} \perp\!\!\!\perp \mathbf{Z}^{T+1} \mid (\{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z})$, for any $\mathbb{Q}(\hat{\mathbf{A}}^t)$ and $\mathbb{Q}(\mathbf{Z}^t)$:

$$I(\mathcal{D}; \mathbf{Z}^{T+1}) \leq I(\mathcal{D}; \{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z}) \leq \sum_{t \in \mathcal{I}_A} \mathcal{A}^t + \sum_{t \in \mathcal{I}_Z} \mathcal{Z}^t, \quad (10)$$

$$\text{where } \mathcal{A}^t = \mathbb{E} \left[\log \frac{\mathbb{P}(\hat{\mathbf{A}}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \mathbf{A}^t)}{\mathbb{Q}(\hat{\mathbf{A}}^t)} \right], \quad (11)$$

$$\mathcal{Z}^t = \mathbb{E} \left[\log \frac{\mathbb{P}(\mathbf{Z}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \hat{\mathbf{A}}^t)}{\mathbb{Q}(\mathbf{Z}^t)} \right]. \quad (12)$$

Proofs for Proposition 1 and 2 are provided in Appendix B.

4.2.2 Deriving the DGIB_C. To ensure \mathbf{Z}^{T+1} adheres to the *Consensual* Constraint with respect to $\mathbf{Z}^{1:T}$, we further constrain the search space Ω as $\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathbf{Z}^{1:T}, C(\theta))$ and optimizing DGIB_C:

$$\begin{aligned} \mathbf{Z}^{T+1} &= \arg \min_{\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathbf{Z}^{1:T}, C(\theta)) \in \Omega} \text{DGIB}_C(\mathbf{Z}^{1:T}, \mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) \\ &\triangleq \left[- \underbrace{I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1})}_{\text{Eq. (15)}} + \beta_2 \underbrace{I(\mathbf{Z}^{1:T}; \mathbf{Z}^{T+1})}_{\text{Eq. (20)}} \right]. \end{aligned} \quad (13)$$

Variational Bounds of DGIB_C. The lower bound of $I(\mathcal{D}; \mathbf{Z}^{T+1})$ in Eq. (13) is consistent with Proposition 1. We introduce the upper bound of $I(\mathbf{Z}^{1:T}; \mathbf{Z}^{T+1})$, which is proved in Appendix B.

Proposition 3 (Upper Bound of $I(\mathbf{Z}^{1:T}; \mathbf{Z}^{T+1})$).

$$I(\mathbf{Z}^{1:T}; \mathbf{Z}^{T+1}) \leq \mathbb{E} \left[\log \frac{\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathbf{Z}^{1:T})}{\mathbb{Q}(\mathbf{Z}^{T+1})} \right]. \quad (14)$$

The distinctions between DGIB_{MS} and DGIB_C primarily pertain to their *inputs* and *search space*. DGIB_{MS} utilizes the original training data $\mathcal{D} = (\mathcal{G}^{1:T}, \mathbf{X}^{T+1})$ with $\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathcal{D}, \theta)$ constraint as input, whereas DGIB_C takes intermediate variables $\mathbf{Z}^{1:T}$ as input with $\mathbb{P}(\mathbf{Z}^{T+1} \mid \mathbf{Z}^{1:T}, C(\theta))$ constraint. DGIB_{MS} and DGIB_C mutually constrains each other via the optimization process, ultimately leading to the satisfaction of the *MSC* Condition approved by \mathbf{Z}^{T+1} .

4.3 DGIB Principle Instantiation

To jointly optimize DGIB_{MS} and DGIB_C, we begin by specifying the lower and upper bounds defined in Proposition 1, 2 and 3.

Instantiation for Eq. (9). To specify the lower bound, we set $\mathbb{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1}) = \text{Cat}(g(\mathbf{Z}^{T+1}))$, where $\text{Cat}(\cdot)$ denotes the categorical distribution, and set $\mathbb{Q}_2(\mathbf{Y}^{T+1}) = \mathbb{P}(\mathbf{Y}^{T+1})$. As the second term in Eq. (9) empirically converges to 1, we ignore it. Thus, the RHS of Eq. (9) reduces to the cross-entropy loss [31], *i.e.*:

$$I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) \doteq -\mathcal{L}_{\text{CE}}(g(\mathbf{Z}^{T+1}), \mathbf{Y}^{T+1}). \quad (15)$$

Instantiation for Eq. (11). We instantiate DGIB on the backbone of the GAT [44], where the attentions can be utilized to refine the initial graph structure or as the parameters of the neighbor sampling distributions. Concretely, let $\phi_{v,k}^t$ as the logits of the attention weights between node v and its spatio-temporal neighbors. We assume the prior of $\mathbb{P}(\hat{\mathbf{A}}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \mathbf{A}^t)$ follows the Bernoulli distribution $\text{Bern}(\cdot)$ or the categorical distribution $\text{Cat}(\cdot)$ both parameterized by $\phi_{v,k}^t$. We assume $\mathbb{Q}(\hat{\mathbf{A}}^t)$ respectively follows the non-informative $\text{Bern}(\cdot)$ or $\text{Cat}(\cdot)$ parameterized by certain constants. Thus, the RHS of Eq. (11) is estimated by:

$$\mathcal{A}^t \doteq \mathbb{E}_{\mathbb{P}(\hat{\mathbf{A}}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \mathbf{A}^t)} \left[\log \frac{\mathbb{P}(\hat{\mathbf{A}}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \mathbf{A}^t)}{\mathbb{Q}(\hat{\mathbf{A}}^t)} \right], \quad (16)$$

which can be further instantiated as:

$$\mathcal{A}_B^t \doteq \sum_{v \in \mathcal{V}^t} \text{KL} \left[\text{Bern}(\phi_{v,k}^t) \parallel \text{Bern}(|\mathcal{N}_{ST}(v, k, t)|^{-1}) \right], \quad (17)$$

$$\text{or } \mathcal{A}_C^t \doteq \sum_{v \in \mathcal{V}^t} \text{KL} \left[\text{Cat}(\phi_{v,k}^t) \parallel \text{Cat}(|\mathcal{N}_{ST}(v, k, t)|^{-1}) \right]. \quad (18)$$

Instantiation for Eq. (12) and Eq. (14). To estimate \mathcal{Z}^t , we set the prior distribution of $\mathbb{P}(\mathbf{Z}^t \mid \hat{\mathbf{Z}}^t, \mathbf{Z}^{t-1}, \hat{\mathbf{A}}^t)$ follows a multivariate

normal distribution $N(\mathbf{Z}^t; \boldsymbol{\mu}_{\mathbb{P}}, \boldsymbol{\sigma}_{\mathbb{P}}^2)$, while $\mathbb{Q}(\mathbf{Z}^t) \sim N(\mathbf{Z}^t; \boldsymbol{\mu}_{\mathbb{Q}}, \boldsymbol{\sigma}_{\mathbb{Q}}^2)$. Inspired by the Markov Chain Monte Carlo (MCMC) sampling [12], Eq. (12) can be estimated with the sampled node set $\mathcal{V}^{t'}$:

$$\mathcal{Z}^t \doteq \sum_{v \in \mathcal{V}^{t'}} \left[\log \Phi(\mathbf{Z}_v^t; \boldsymbol{\mu}_{\mathbb{P}}, \boldsymbol{\sigma}_{\mathbb{P}}^2) - \log \Phi(\mathbf{Z}_v^t; \boldsymbol{\mu}_{\mathbb{Q}}, \boldsymbol{\sigma}_{\mathbb{Q}}^2) \right], \quad (19)$$

where $\Phi(\cdot)$ is the Probability Density Function (PDF) of the normal distribution. Similarly, we specify Eq. (14) with sampled $\mathcal{V}^{(T+1)'}$:

$$\sum_{v \in \mathcal{V}^{(T+1)'}} \left[\log \Phi(\mathbf{Z}_v^{T+1}; \boldsymbol{\mu}_{\mathbb{P}}, \boldsymbol{\sigma}_{\mathbb{P}}^2) - \log \Phi(\mathbf{Z}_v^{T+1}; \boldsymbol{\mu}_{\mathbb{Q}}, \boldsymbol{\sigma}_{\mathbb{Q}}^2) \right]. \quad (20)$$

Training Objectives. To acquire a tractable version of Eq. (6), we first plug Eq. (15), Eq. (17)/Eq. (18) and Eq. (19) into Eq. (8), respectively, to estimate DGIB_{MS} . Then plug Eq. (15) and Eq. (20) into Eq. (13) to estimate DGIB_C . The overall training objectives of the proposed DGIB can be rewritten as:

$$\mathcal{L}_{\text{DGIB}} = \alpha \text{DGIB}_{MS} + (1 - \alpha) \text{DGIB}_C, \quad (21)$$

where the α is a trade-off hyperparameter.

4.4 Optimization and Complexity Analysis

The overall training pipeline of DGIB is shown in Algorithm 1. With the proposed upper and lower bounds for intractable terms in both DGIB_{MS} and DGIB_C , the overall framework can be trained end-to-end using back-propagation, and thus we can use gradient descent to optimize. Based on the detailed analysis in Appendix A, the time complexity of our method is:

$$O(2Lk(T+1)|\mathcal{V}||\mathcal{E}|dd'^2) + O(|\mathcal{E}'|d'), \quad (22)$$

which is on par with the state-of-the-art DGNNs. We further illustrate the training time efficiency in Appendix C.4.

5 EXPERIMENT

In this section, we conduct extensive experiments on both real-world and synthetic dynamic graph datasets to evaluate the robustness of our DGIB against adversarial attacks. We first introduce the experimental settings and then present the results. Additional configurations and results can be found in Appendix C and D.

5.1 Experimental Settings

5.1.1 Dynamic Graph Datasets. In order to comprehensively evaluate the effectiveness of our proposed method, we use three real-world dynamic graph datasets to evaluate DGIB^1 on the challenging future link prediction task. **COLLAB** [41] is an academic collaboration dataset with papers published in 16 years, which reveals the dynamic citation networks among authors. **Yelp** [33] contains customer reviews on business for 24 months, which are collected from the crowd-sourced local business review and social networking web. **ACT** [22] describes the actions taken by users on a popular MOOC website within 30 days, and each action has a binary label. Statistics of the datasets are concluded in Table B.1, which also contains the split of snapshots for training, validation, and testing.

¹The code of DGIB is available at <https://github.com/RingBDStack/DGIB> and <https://doi.org/10.5281/zenodo.10633302>.

5.1.2 Baselines. We compare DGIB with three categories.

- **Static GNNs: GAE** [20] and **VGAE** [20] are representative GCN [19]-based autoencoders. **GAT** [44] using attention mechanisms to dynamically aggregate node features, which is also the default backbone of DGIB.
- **Dynamic GNNs: GCRN** [35] adopts GCNs to obtain node embeddings, followed by a GRU [8] to capture temporal relations. **EvolveGCN** [30] applies an LSTM [15] or GRU [8] to evolve the parameters of GCNs. **DySAT** [33] models self-attentions in both structural and temporal domains.
- **Robust and Generalized (D)GNNs: IRM** [2] learns robust representation by minimizing invariant risk. **V-REx** [21] extends IRM [2] by reweighting the risk. **GroupDRO** [32] reduces the risk gap over training distributions. **RGCN** [58] fortifies GCNs against adversarial attacks by Gaussian reparameterization and variance-based attention. **DIDA** [56] exploits robust and generalized predictive patterns on dynamic graphs. **GIB** [46] is the most relevant baseline to ours, which learns robust representations with structure-involved IB principle.

5.1.3 Adversarial Attack Settings. We compare baselines and the proposed DGIB under two adversarial attack settings.

- **Non-targeted Adversarial Attack:** We produce synthetic datasets by attacking *graph structures* and *node features*, respectively. **(1) Attack graph structures.** We randomly remove one out of five types of links in training and validation graphs in each dataset (information on link type has been removed after the above operations), which is more practical and challenging in real-world scenarios as the model cannot get access to any features about the filtered links. **(2) Attack node features.** We add random Gaussian noise $\lambda \cdot r \cdot \epsilon$ to each dimension of the node features for all nodes, where r is the reference amplitude of original features, and $\epsilon \sim N(\mathbf{0}, \mathbf{I})$. $\lambda \in \{0.5, 1.0, 1.5\}$ acts as the parameter to control the attacking degree.
- **Targeted Adversarial Attack:** We apply the prevailing **NETTACK** [59], a strong targeted adversarial attack library on graphs designed to target nodes by altering their connected edges or node features. We simultaneously consider the *evasion* attack and *poisoning* attack. **(1) Evasion attack.** We train the model on clean datasets and perform attacking on each graph snapshot in the testing split. **(2) Poisoning attack.** We attack the whole dataset before model training and testing. In both scenarios, we follow the default settings of **NETTACK** [59] to select targeted attacking nodes, and choose **GAT** [44] as the surrogate model with default parameter settings. We set the number of perturbations n in $\{1, 2, 3, 4\}$.

5.1.4 Hyperparameter Settings. We set the number of layers as two for baselines, and as one for DGIB to avoid overfitting. We set the representation dimension of all baselines and our DGIB to be 128. The hyperparameters of baselines are set as the suggested value in their papers or carefully tuned for fairness. The suggested values of α , β_1 and β_2 can be found in the configuration files. For the optimization, we use Adam [18] with a learning rate selected from $\{1e-03, 1e-04, 1e-05, 1e-06\}$ adopt the grid search for the best performance using the validation split. We set the maximum epoch number as 1000 with the early stopping strategy.

Table 1: AUC score ($\% \pm$ standard deviation) of future link prediction task on real-world datasets against *non-targeted* adversarial attacks. The best results are shown in bold type and the runner-ups are underlined.

Dataset	COLLAB					Yelp					ACT				
	Clean	Structure Attack	Feature Attack			Clean	Structure Attack	Feature Attack			Clean	Structure Attack	Feature Attack		
Model			$\lambda = 0.5$	$\lambda = 1.0$	$\lambda = 1.5$			$\lambda = 0.5$	$\lambda = 1.0$	$\lambda = 1.5$			$\lambda = 0.5$	$\lambda = 1.0$	$\lambda = 1.5$
GAE [20]	77.15 \pm 0.5	74.04 \pm 0.8	50.59 \pm 0.8	44.66 \pm 0.8	43.12 \pm 0.8	70.67 \pm 1.1	64.45 \pm 5.0	51.05 \pm 0.6	45.41 \pm 0.6	41.56 \pm 0.9	72.31 \pm 0.5	60.27 \pm 0.4	56.56 \pm 0.5	52.52 \pm 0.6	50.36 \pm 0.9
VGAE [20]	86.47 \pm 0.0	74.95 \pm 1.2	56.75 \pm 0.6	50.39 \pm 0.7	48.68 \pm 0.7	76.54 \pm 0.5	65.33 \pm 1.4	55.53 \pm 0.7	49.88 \pm 0.8	45.08 \pm 0.6	79.18 \pm 0.5	66.29 \pm 1.3	60.67 \pm 0.7	57.39 \pm 0.8	55.27 \pm 1.0
GAT [44]	88.26 \pm 0.4	77.29 \pm 1.8	58.13 \pm 0.9	51.41 \pm 0.9	49.77 \pm 0.9	77.93 \pm 0.1	69.35 \pm 1.6	56.72 \pm 0.3	52.51 \pm 0.5	46.21 \pm 0.5	85.07 \pm 0.3	77.55 \pm 1.2	66.05 \pm 0.4	61.85 \pm 0.3	59.05 \pm 0.3
GCRN [35]	82.78 \pm 0.5	69.72 \pm 0.5	54.07 \pm 0.9	47.78 \pm 0.8	46.18 \pm 0.9	68.59 \pm 1.0	54.68 \pm 7.6	52.68 \pm 0.6	46.85 \pm 0.6	40.45 \pm 0.6	76.28 \pm 0.5	64.35 \pm 1.2	59.48 \pm 0.7	54.16 \pm 0.6	53.88 \pm 0.7
EvolveGCN [30]	86.62 \pm 1.0	76.15 \pm 0.9	56.82 \pm 1.2	50.33 \pm 1.0	48.55 \pm 1.0	78.21 \pm 0.0	53.82 \pm 2.0	57.91 \pm 0.5	51.82 \pm 0.3	45.32 \pm 1.0	74.55 \pm 0.3	63.17 \pm 1.0	61.02 \pm 0.5	53.34 \pm 0.5	51.62 \pm 0.7
DySAT [33]	88.77 \pm 0.2	76.59 \pm 0.2	58.28 \pm 0.3	51.52 \pm 0.3	49.32 \pm 0.5	78.87 \pm 0.6	66.09 \pm 1.4	58.46 \pm 0.4	52.33 \pm 0.7	46.24 \pm 0.7	78.52 \pm 0.4	66.55 \pm 1.2	61.94 \pm 0.8	56.98 \pm 0.8	54.14 \pm 0.7
IRM [2]	87.96 \pm 0.9	75.42 \pm 0.9	60.51 \pm 1.3	53.89 \pm 1.1	52.17 \pm 0.9	66.49 \pm 10.8	56.02 \pm 16.0	50.96 \pm 3.3	48.58 \pm 5.2	45.32 \pm 3.3	80.02 \pm 0.6	69.19 \pm 1.4	62.84 \pm 0.1	57.28 \pm 0.2	56.04 \pm 0.2
V-REx [20]	88.31 \pm 0.3	76.24 \pm 0.8	61.23 \pm 1.5	54.51 \pm 1.0	52.24 \pm 1.1	79.04 \pm 0.2	66.41 \pm 1.9	61.49 \pm 0.5	53.72 \pm 1.0	51.32 \pm 0.9	83.11 \pm 0.3	70.15 \pm 1.1	65.59 \pm 0.1	60.03 \pm 0.3	58.79 \pm 0.2
GroupDRO [32]	88.76 \pm 0.1	76.33 \pm 0.3	61.10 \pm 1.3	54.62 \pm 1.0	52.33 \pm 0.8	<u>79.38\pm0.4</u>	66.97 \pm 0.6	61.78 \pm 0.8	55.37 \pm 0.9	52.18 \pm 0.7	85.19 \pm 0.5	74.35 \pm 1.6	66.05 \pm 0.5	61.85 \pm 0.4	59.05 \pm 0.3
RGCN [58]	88.21 \pm 0.1	78.66 \pm 0.7	61.29 \pm 0.5	54.29 \pm 0.6	52.99 \pm 0.6	<u>77.28\pm0.3</u>	74.29 \pm 0.4	59.72 \pm 0.3	52.88 \pm 0.3	50.40 \pm 0.2	87.22 \pm 0.2	82.66 \pm 0.4	68.51 \pm 0.2	62.67 \pm 0.2	61.31 \pm 0.2
DIDA [56]	91.97 \pm 0.0	80.87 \pm 0.4	61.32 \pm 0.8	55.77 \pm 0.9	54.91 \pm 0.9	78.22 \pm 0.4	<u>75.92\pm0.9</u>	60.83 \pm 0.6	54.11 \pm 0.6	50.21 \pm 0.6	89.84 \pm 0.8	78.64 \pm 1.0	70.97 \pm 0.2	64.49 \pm 0.4	62.57 \pm 0.2
GIB [46]	91.36 \pm 0.2	80.89 \pm 0.1	61.88 \pm 0.8	55.15 \pm 0.8	54.65 \pm 0.9	77.52 \pm 0.4	<u>75.03\pm0.3</u>	<u>61.94\pm0.9</u>	<u>56.15\pm0.3</u>	<u>52.21\pm0.8</u>	92.33 \pm 0.3	86.99 \pm 0.3	72.16 \pm 0.5	66.72 \pm 0.2	64.96 \pm 0.5
DGIB-Bern	<u>92.17\pm0.2</u>	<u>83.58\pm0.1</u>	<u>63.54\pm0.9</u>	<u>56.92\pm1.0</u>	<u>56.24\pm1.0</u>	76.88 \pm 0.2	75.61 \pm 0.0	63.91\pm0.9	59.28\pm0.9	54.77\pm1.0	<u>94.49\pm0.2</u>	<u>87.75\pm0.1</u>	<u>73.05\pm0.9</u>	<u>68.49\pm0.9</u>	66.27\pm0.9
DGIB-Cat	92.68\pm0.1	84.16\pm0.1	63.99\pm0.5	57.76\pm0.8	<u>55.63\pm1.0</u>	<u>79.53\pm0.2</u>	<u>77.72\pm0.1</u>	61.42 \pm 0.9	55.12 \pm 0.7	51.90 \pm 0.9	94.89\pm0.2	88.27\pm0.2	73.92\pm0.8	68.88\pm0.9	<u>65.99\pm0.7</u>

5.2 Against Non-targeted Adversarial Attacks

In this section, we evaluate model performance on the future link prediction task, as well as the robustness against non-targeted adversarial attacks in terms of graph structures and node features. Specifically, we train baselines and our DGIB on the clean datasets, after which we perturb edges and features respectively on the testing split following the experimental settings. Note that, DGIB with a prior of the Bernoulli distribution as Eq. (17) is referred to as DGIB-Bern, while DGIB with a prior of the categorical distribution in Eq. (18) is denoted as DGIB-Cat. Results are reported with the metric of AUC (%) score in five runs and concluded in Table 1.

Results. DGIB-Bern and DGIB-Cat outperform GIB [46] and other baselines in most scenarios. Static GNN baselines fail in all settings as they cannot model dynamics. Dynamic GNN baselines fail due to their weak robustness. Some robust and generalized (D)GNNs, such as GroupDRO [32], DIDA [56] and GIB [46] outperform DGIB-Bern or DGIB-Cat slightly in a few scenarios, but they generally fall behind DGIB in most cases due to the insufficient feature compression and conservation, which greatly impact the model robustness against the non-targeted adversarial attacks.

5.3 Against Targeted Adversarial Attacks

In this section, we continue to compare the proposed DGIB with competitive baselines standing out in Table 1, considering the link prediction performance and robustness against targeted adversarial attacks, which reveals whether we successfully defended the attacks. Specifically, we generate attacked datasets with NETTACK [59] concerning different perturbation times n in both evasion attacking mode and poisoning attacking mode. Higher n represents a heavier attacking degree. Results are reported in Table 2.

Results. DGIB-Bern and DGIB-Cat outperform all baselines under different settings. Results demonstrate DGIB is contained under challenging evasion and poisoning adversarial attacks. Comprehensively, DGIB-Bern owns a better performance in targeted adversarial attacks with a lower average AUC decrease, while DGIB-Bern is better in non-targeted adversarial attacks. We explain this phenomenon as the $\text{Cat}(\cdot)$ is the non-informative distribution, which

fits well to the non-targeted settings, and $\text{Bern}(\cdot)$ requires priors, which may be more appropriate for against targeted attacks.

5.4 Ablation Study

In this section, we analyze the effectiveness of the three variants:

- **DGIB (w/o Cons):** We remove the *Consensual* channel DGIB_C in the overall training objective (Eq. (21)), and optimizing only with the *Minimal* and *Sufficient* channel DGIB_{MS} .
- **DGIB (w/o \mathcal{A}):** We remove the structure sampling term (\mathcal{A}) in the upper bound of $I(\mathcal{D}; \mathbf{Z}^{T+1})$ (Eq. (10)).
- **DGIB (w/o \mathcal{Z}):** We remove the feature sampling term (\mathcal{Z}) in the upper bound of $I(\mathcal{D}; \mathbf{Z}^{T+1})$ (Eq. (10)).

We choose DGIB-Bern as the backbone and compare performances on the clean, evasion attacked ($n = 2$) and poisoning attacked ($n = 2$) COLLAB, respectively. Results are shown in Figure 4.

Results. Overall, DGIB outperforms the other three variants, except for DGIB (w/o \mathcal{A}), where it exceeds the original DGIB-Bern on the clean COLLAB by 1.1%. We claim this phenomenon is within our expectation as the structure sampling term (\mathcal{A}) contributes to raising the robustness by refining structures and compression feature information, which will surely damage its performance on the clean dataset. Concretely, we witness DGIB surpassing all three variants when confronting evasion and poisoning adversarial attacks, which provides insights into the effectiveness of the proposed components and demonstrates their importance in achieving better performance for robust representation learning on dynamic graphs.

5.5 Information Plane Analysis

In this section, we observe the evolution of the IB compression process on the Information Plane, which is widely applied to analyze the changes in the mutual information between input, latent representations, and output during training. Given the Markov Chain $\langle \mathbf{X} \rightarrow \mathbf{Y} \rightarrow \mathbf{Z} \rangle$, the latent representation is uniquely mapped to a point in the Information Plane with coordinates $(I(\mathbf{X}; \mathbf{Z}), I(\mathbf{Y}; \mathbf{Z}))$. We analyze DGIB-Bern on the clean COLLAB and draw the coordinates of $(I(\mathcal{D}; \mathbf{Z}^{T+1}), I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}))$ in Figure 5.

Table 2: AUC score (% ± standard deviation) of future link prediction task on real-world datasets against *targeted* adversarial attacks. The best results are shown in bold type and the runner-ups are underlined.

Dataset	Model	Clean	Evasion Attack					Poisoning Attack				
			$n = 1$	$n = 2$	$n = 3$	$n = 4$	Avg. Decrease	$n = 1$	$n = 2$	$n = 3$	$n = 4$	Avg. Decrease
COLLAB	VGAE [20]	86.47±0.0	73.39±0.1	62.18±0.1	51.72±0.1	46.97±0.1	↓ 32.27	63.42±0.3	52.63±0.3	50.98±0.4	45.64±0.3	↓ 38.51
	GAT [44]	88.26±0.4	76.21±0.1	66.56±0.1	57.92±0.1	50.96±0.1	↓ 28.71	66.59±0.5	55.31±0.6	51.34±0.7	48.99±0.9	↓ 37.05
	DySAT [33]	88.77±0.2	77.91±0.1	68.22±0.1	58.82±0.1	51.39±0.1	↓ 27.80	69.02±0.3	57.62±0.3	52.76±0.3	50.07±0.8	↓ 35.37
	RGCN [58]	88.21±0.1	77.65±0.1	67.11±0.1	59.06±0.1	52.02±0.1	↓ 27.49	69.48±0.2	58.39±0.3	52.48±0.6	50.62±0.9	↓ 34.53
	GIB [46]	91.36±0.2	78.95±0.0	69.63±0.1	60.98±0.0	54.48±0.2	↓ 27.74	71.47±0.3	<u>61.03±0.4</u>	54.97±0.7	52.09±1.0	↓ 34.44
	DGIB-Bern	<u>92.17±0.2</u>	81.36±0.0	72.79±0.0	63.25±0.1	57.22±0.1	↓ 25.51	74.06±0.3	61.93±0.2	56.57±0.2	<u>52.62±0.3</u>	↓ 33.49
	DGIB-Cat	92.68±0.1	<u>81.29±0.0</u>	<u>71.32±0.1</u>	<u>62.03±0.1</u>	<u>55.08±0.1</u>	↓ 27.24	<u>72.55±0.2</u>	<u>60.99±0.3</u>	<u>55.62±0.4</u>	53.08±0.3	↓ 34.65
Yelp	VGAE [20]	76.54±0.5	65.86±0.1	54.82±0.2	48.08±0.1	46.25±0.1	↓ 29.77	62.73±0.6	52.61±0.4	47.72±0.4	45.43±0.5	↓ 31.90
	GAT [44]	77.93±0.1	67.96±0.1	59.47±0.1	50.27±0.1	48.62±0.1	↓ 27.39	65.34±0.5	54.51±0.2	50.24±0.4	48.96±0.4	↓ 29.72
	DySAT [33]	78.87±0.6	69.77±0.1	60.66±0.1	52.16±0.1	50.15±0.1	↓ 26.22	66.87±0.6	56.31±0.3	50.44±0.6	<u>50.49±0.5</u>	↓ 28.96
	RGCN [58]	77.28±0.3	68.54±0.1	60.69±0.1	51.51±0.1	49.72±0.1	↓ 25.44	65.55±0.4	55.47±0.3	49.08±0.6	49.09±0.6	↓ 29.09
	GIB [46]	77.52±0.4	68.59±0.1	<u>61.22±0.1</u>	51.26±0.1	49.58±0.1	↓ 25.61	65.59±0.3	<u>56.79±0.3</u>	50.92±0.4	49.55±0.4	↓ 28.13
	DGIB-Bern	76.88±0.2	72.27±0.1	60.96±0.0	54.32±0.1	51.73±0.1	↓ 22.19	68.64±0.2	56.73±0.2	53.18±0.3	50.21±0.2	↓ 25.61
	DGIB-Cat	79.53±0.2	<u>70.17±0.0</u>	62.25±0.1	<u>52.69±0.1</u>	<u>50.87±0.1</u>	↓ 25.82	<u>67.38±0.3</u>	57.02±0.2	<u>51.39±0.2</u>	50.53±0.2	↓ 28.85
ACT	VGAE [20]	79.18±0.5	67.59±0.1	62.98±0.1	54.33±0.1	52.26±0.0	↓ 25.11	62.55±1.6	55.15±1.7	51.02±1.8	50.11±1.9	↓ 30.90
	GAT [44]	85.07±0.3	75.14±0.1	67.25±0.1	59.75±0.1	58.51±0.1	↓ 23.40	71.26±0.9	61.43±1.1	57.35±1.1	58.53±1.0	↓ 26.95
	DySAT [33]	78.52±0.4	70.64±0.1	63.35±0.0	56.36±0.0	55.12±0.1	↓ 21.84	66.21±0.9	56.28±0.9	53.45±1.1	54.43±1.0	↓ 26.65
	RGCN [58]	87.22±0.2	78.64±0.1	70.11±0.1	62.99±0.1	61.31±0.1	↓ 21.73	73.71±0.8	63.43±0.9	59.97±1.3	60.41±0.8	↓ 26.18
	GIB [46]	92.33±0.3	<u>85.61±0.1</u>	74.08±0.1	65.44±0.1	64.04±0.1	↓ 21.70	80.01±0.7	67.04±0.8	63.85±0.6	60.95±0.7	↓ 26.39
	DGIB-Bern	<u>94.49±0.2</u>	89.83±0.1	85.81±0.1	79.95±0.1	78.01±0.1	↓ 11.73	80.92±0.3	70.76±0.4	65.27±0.6	61.93±0.9	↓ 26.21
	DGIB-Cat	94.89±0.2	84.98±0.1	<u>76.78±0.1</u>	<u>67.69±0.1</u>	<u>66.68±0.1</u>	↓ 21.98	80.16±0.4	<u>68.71±0.5</u>	<u>64.38±0.6</u>	65.43±0.9	↓ 26.57




Figure 4: Ablation study results.

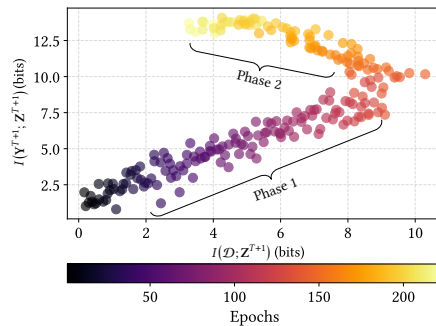


Figure 5: Information Plane analysis.

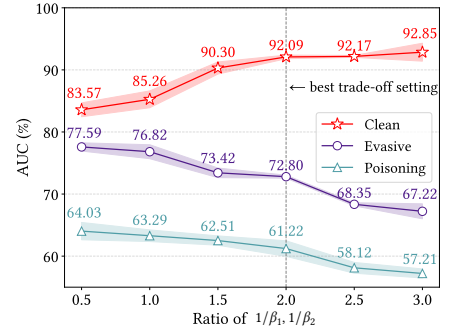


Figure 6: Performance trade-off analysis.

Results. The information evolution process is composed of two phases. During Phase 1 (ERM phase), $I(\mathcal{D}; Z^{T+1})$ and $I(Y^{T+1}; Z^{T+1})$ both increase, indicating the latent representations are extracting information about the input and labels, and the coordinates are shifting toward the upper right corner. In Phase 2 (compression phase), $I(\mathcal{D}; Z^{T+1})$ begins to decline, while the growth rate of $I(Y^{T+1}; Z^{T+1})$ slows down, and converging to the upper left corner, indicating our DGIB begins to take effect, leading to a *Minimal* and *Sufficient* latent representation with *Consensual Constraint*.

5.6 Hyperparameter Trade-off Analysis

In this section, we analyze the impact of the compression parameters β_1 and β_2 on the trade-off between the performance of prediction and robustness. We conduct experiments based on DGIB-Bern with different ratios of MSE term and compression term ($1/\beta_1, 1/\beta_2$) on the clean, evasion attacked ($n = 2$) and poisoning attacked ($n = 2$) COLLAB, respectively. Results are reported in Figure 6.

Results. With the increase of $1/\beta_1$ or $1/\beta_2$, DGIB-Bern performs better on the clean COLLAB, while its robustness against targeted adversarial attacks decreases. This validates the conflicts between the MSE term and the compression term, as the compression term

sacrifices prediction performance to improve the robustness. We find the optimal setting for $1/\beta_1$ and $1/\beta_2$, where both the performance of prediction and robustness are well-preserved.

6 CONCLUSION

In this paper, we present a novel framework named DGIB to learn robust and discriminative representations on dynamic graphs grounded in the information-theoretic IB principle for the first time. We decompose the overall DGIB objectives into $DGIB_M$ and $DGIB_C$ channels, which act jointly to satisfy the proposed MSC Condition that the optimal representations should satisfy. Variational bounds are further introduced to efficiently and appropriately estimate intractable IB terms. Extensive experiments demonstrate the superior robustness of DGIB against adversarial attacks compared with state-of-the-art baselines in the link prediction task.

ACKNOWLEDGMENTS

The corresponding author is Qingyun Sun. The authors of this paper are supported by the National Natural Science Foundation of China through grants No.62225202, and No.62302023. We owe sincere thanks to all authors for their valuable efforts and contributions.

REFERENCES

- [1] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2017. Deep Variational Information Bottleneck. In *ICLR*.
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant Risk Minimization. *arXiv* (2019).
- [3] Normand J Beaudry and Renato Renner. 2012. An Intuitive Proof of the Data Processing Inequality. *Quantum Information & Computation* 12 (2012), 432–441.
- [4] Tanya Y Berger-Wolf and Jared Saia. 2006. A Framework for Analysis of Dynamic Social Networks. In *KDD*. 523–528.
- [5] Lei Cai, Zhengzhang Chen, Chen Luo, Jiaping Gui, Jingchao Ni, Ding Li, and Haifeng Chen. 2021. Structural Temporal Graph Neural Networks for Anomaly Detection in Dynamic Graphs. In *CIKM*. 3747–3756.
- [6] Zhiyuan Cai, Kaiqi Zhao, Kenny Q Zhu, and Haixun Wang. 2013. Wikification via Link Co-occurrence. In *CIKM*. 1087–1096.
- [7] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. 2020. Measuring and Relieving the Over-smoothing Problem for Graph Neural Networks from the Topological View. In *AAAI*, Vol. 34. 3438–3445.
- [8] Kyunghyun Cho, Bart van Merriënboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations Using RNN Encoder-decoder for Statistical Machine Translation. In *EMNLP*. 1724–1734.
- [9] Jiameng Fan and Wenchao Li. 2022. DRIBO: Robust Deep Reinforcement Learning via Multi-View Information Bottleneck. In *ICML*. 6074–6102.
- [10] Dongqi Fu and Jingrui He. 2021. SDG: A Simplified and Dynamic Graph Neural Network. In *SIGIR*. 2273–2277.
- [11] Xingcheng Fu, Yuecen Wei, Qingyun Sun, Haonan Yuan, Jia Wu, Hao Peng, and Jianxin Li. 2023. Hyperbolic Geometric Graph Representation Learning for Hierarchy-imbalance Node Classification. In *WWW*. 460–468.
- [12] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. 1995. *Markov Chain Monte Carlo in Practice*.
- [13] Gordon, Greenspan, and Goldberger. 2003. Applying the Information Bottleneck Principle to Unsupervised Clustering of Discrete and Continuous Image Representations. In *ICCV*. 370–377.
- [14] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. 2021. Dynamic Neural Networks: A Survey. *IEEE TPAMI* 44, 11 (2021), 7436–7456.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-term Memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [16] Maximilian Igl, Kamil Ciosek, Yingzhen Li, Sebastian Tschitschek, Cheng Zhang, Sam Devlin, and Katja Hofmann. 2019. Generalization in Reinforcement Learning with Selective Noise Injection and Information Bottleneck. In *NeurIPS*, Vol. 32.
- [17] Masahiro Ito, Kotaro Nakayama, Takahiro Hara, and Shojiro Nishio. 2008. Association Thesaurus Construction Methods Based on Link Co-occurrence Analysis for Wikipedia. In *CIKM*. 817–826.
- [18] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [19] Thomas N Kipf and Max Welling. 2016. Semi-supervised Classification with Graph Convolutional Networks. *arXiv* (2016).
- [20] Thomas N Kipf and Max Welling. 2016. Variational Graph Auto-encoders. *arXiv* (2016).
- [21] David Krueger, Ethan Caballero, Jörn-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Rémi Le Priol, and Aaron C. Courville. 2021. Out-of-distribution Generalization via Risk Extrapolation. In *ICML*, Vol. 139. 5815–5826.
- [22] Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks. In *KDD*. 1269–1278.
- [23] Honglin Li, Chenglu Zhu, Yunlong Zhang, Yuxuan Sun, Zhongyi Shui, Wenwei Kuang, Sunyi Zheng, and Lin Yang. 2023. Task-Specific Fine-Tuning via Variational Information Bottleneck for Weakly-supervised Pathology Whole Slide Image Classification. In *CVPR*. 7454–7463.
- [24] Mengzhang Li and Zhanxing Zhu. 2021. Spatial-temporal Fusion Graph Neural Networks for Traffic Flow Forecasting. In *AAAI*, Vol. 35. 4189–4196.
- [25] Antonio Longa, Steve Azzolin, Gabriele Santin, Giulia Cencetti, Pietro Liò, Bruno Lepri, and Andrea Passerini. 2022. Explaining the Explainers in Graph Neural Networks: A Comparative Study. *arXiv* (2022).
- [26] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv* (2013).
- [27] Vinod Nair and Geoffrey E Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *ICML*. 807–814.
- [28] XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. 2010. Estimating Divergence Functionals and the Likelihood Ratio by Convex Risk Minimization. *IEEE Transactions on Information Theory* 56, 11 (2010), 5847–5861.
- [29] Bhargavi Paranjape, Mandar Joshi, John Thickstun, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. An Information Bottleneck Approach for Controlling Conciseness in Rationale Extraction. In *EMNLP*. 1938–1952.
- [30] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. 2020. EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs. In *AAAI*, Vol. 34. 5363–5370.
- [31] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. 2019. On Variational Bounds of Mutual Information. In *ICML*. 5171–5180.
- [32] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally Robust Neural Networks for Group Shifts: On The Importance of Regularization for Worst-case Generalization. *arXiv* (2019).
- [33] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. 2020. DySAT: Deep Neural Representation Learning on Dynamic Graphs via Self-attention Networks. In *WSDM*. 519–527.
- [34] Michael Schweinberger and Mark S Handcock. 2015. Local Dependence in Random Graph Models: Characterization, Properties and Statistical Inference. *Journal of the Royal Statistical Society Series B* 77, 3 (2015), 647–676.
- [35] Youngjoon Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. 2018. Structured Sequence Modeling with Graph Convolutional Recurrent Networks. In *ICONIP*. 362–373.
- [36] Ravid Shwartz-Ziv and Naftali Tishby. 2017. Opening the Black Box of Deep Neural Networks via Information. *arXiv* (2017).
- [37] Qingyun Sun, Jianxin Li, Hao Peng, Jia Wu, Xingcheng Fu, Cheng Ji, and S Yu Philip. 2022. Graph Structure Learning with Variational Information Bottleneck. In *AAAI*, Vol. 36. 4165–4174.
- [38] Qingyun Sun, Jianxin Li, Hao Peng, Jia Wu, Yuanxing Ning, Philip S Yu, and Lifang He. 2021. SUGAR: Subgraph Neural Network with Reinforcement Pooling and Self-supervised Mutual Information Mechanism. In *WWW*. 2081–2091.
- [39] Qingyun Sun, Jianxin Li, Haonan Yuan, Xingcheng Fu, Hao Peng, Cheng Ji, Qian Li, and Philip S Yu. 2022. Position-aware Structure Learning for Graph Topology-imbalance by Relieving Under-reaching and Over-squashing. In *CIKM*. 1848–1857.
- [40] Hui Tang and Xun Liang. 2023. Where to Find Fascinating Inter-Graph Supervision: Imbalanced Graph Classification with Kernel Information Bottleneck. In *ACM MM*. 3240–3249.
- [41] Jie Tang, Sen Wu, Jimeng Sun, and Hang Su. 2012. Cross-domain Collaboration Recommendation. In *KDD*. 1285–1293.
- [42] Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. The Information Bottleneck Method. *arXiv* (2000).
- [43] Naftali Tishby and Noga Zaslavsky. 2015. Deep Learning and the Information Bottleneck Principle. In *IEEE Information Theory Workshop*. 1–5.
- [44] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.
- [45] Yuecen Wei, Haonan Yuan, Xingcheng Fu, Qingyun Sun, Hao Peng, Xianxian Li, and Chunming Hu. 2024. Poincaré Differential Privacy for Hierarchy-aware Graph Embedding. In *AAAI*.
- [46] Tailin Wu, Hongyu Ren, Pan Li, and Jure Leskovec. 2020. Graph Information Bottleneck. In *NeurIPS*, Vol. 33. 20437–20448.
- [47] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *ICLR*.
- [48] Kuo Yang, Zhengyang Zhou, Wei Sun, Pengkun Wang, Xu Wang, and Yang Wang. 2023. EXTRACT and REFINE: Finding a Support Subgraph Set for Graph Representation. In *KDD*. 2953–2964.
- [49] Liang Yang, Yuanfang Guo, Junhua Gu, Di Jin, Bo Yang, and Xiaochun Cao. 2020. Probabilistic Graph Convolutional Network via Topology-constrained Latent Space Model. *IEEE Transactions on Cybernetics* 52, 4 (2020), 2123–2136.
- [50] Liang Yang, Zesheng Kang, Xiaochun Cao, Di Jin, Bo Yang, and Yuanfang Guo. 2019. Topology Optimization based Graph Convolutional Network. In *IJCAI*. 4054–4061.
- [51] Liang Yang, Fan Wu, Zichen Zheng, Bingxin Niu, Junhua Gu, Chuan Wang, Xiaochun Cao, and Yuanfang Guo. 2021. Heterogeneous Graph Information Bottleneck. In *IJCAI*. 1638–1645.
- [52] Junchi Yu, Tingyang Xu, Yu Rong, Yatao Bian, Junzhou Huang, and Ran He. 2021. Recognizing Predictive Substructures with Subgraph Information Bottleneck. *IEEE TPAMI* (2021).
- [53] Haonan Yuan, Qingyun Sun, Xingcheng Fu, Ziwei Zhang, Cheng Ji, Hao Peng, and Jianxin Li. 2023. Environment-Aware Dynamic Graph Learning for Out-of-Distribution Generalization. In *NeurIPS*.
- [54] Shilei Zhang, Toyotaro Suzumura, and Li Zhang. 2021. DynGraphTrans: Dynamic Graph Embedding via Modified Universal Transformer Networks for Financial Transaction Data. In *SMDS*. 184–191.
- [55] Xiyue Zhang, Chao Huang, Yong Xu, Lianghao Xia, Peng Dai, Liefeng Bo, Junbo Zhang, and Yu Zheng. 2021. Traffic Flow Forecasting with Spatial-temporal Graph Diffusion Network. In *AAAI*, Vol. 35. 15008–15015.
- [56] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Zhou Qin, and Wenwu Zhu. 2022. Dynamic Graph Neural Networks under Spatio-Temporal Distribution Shift. In *NeurIPS*, Vol. 35. 6074–6089.
- [57] Yanping Zheng, Zhewei Wei, and Jiajun Liu. 2023. Decoupled Graph Neural Networks for Large Dynamic Graphs. *VLDB* 16, 9 (2023), 2239–2247.
- [58] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2019. Robust Graph Convolutional Networks Against Adversarial Attacks. In *KDD*. 1399–1407.
- [59] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial Attacks on Neural Networks for Graph Data. In *KDD*. 2847–2856.

A DETAILED OPTIMIZATION AND COMPLEXITY ANALYSIS

We illustrate the overall training pipeline of DGIB in Algorithm 1.

Algorithm 1: Overall training pipeline of DGIB framework.

Input: Dynamic graph $DG = \{\mathcal{G}^t\}_{t=1}^T$; Node features $\mathbf{X}^{1:T+1}$; Labels $\mathbf{Y}^{1:T}$ of link occurrence; Number of layers L ; Relative time encoding function $\text{RTE}(\cdot)$; Non-linear rectifier τ ; Activation function σ ; Hyperparameters k, α, β_1 and β_2 .

Output: The optimized robust model $f_{\theta}^* = w \circ g$; Predicted label $\hat{\mathbf{Y}}^{T+1}$ of link occurrence at time $T + 1$.

- 1 Initialize weights \mathbf{W} and learnable parameters θ randomly;
- 2 $\mathbf{Z}^{1:T+1, (0)} \leftarrow \text{RTE}(\mathbf{X}^{1:T+1})$, attentions $\mathbf{w} \leftarrow \text{GAT}(\mathbf{Z}^{1:T+1, (0)})$;
- 3 Construct $\mathcal{N}_{ST}(v, k, t) \leftarrow \{u \in \mathcal{V}^{t-1:t} \mid d(u, v) = k\}$;
- 4 **for** $l = 1, 2, \dots, L$ **and** $v \in \mathcal{V}^{1:T+1}$ **do**
- 5 **for** t in range $[T + 1]$ **do**
- 6 $\hat{\mathbf{Z}}^{t(l-1)} \leftarrow \tau(\mathbf{Z}^{t(l-1)})\mathbf{W}^{(l)}$;
- 7 $\phi_{v,k}^{t(l)} \leftarrow \sigma\{\langle \hat{\mathbf{Z}}_v^{t(l-1)} \parallel \hat{\mathbf{Z}}_u^{t-1:t(l-1)} \rangle \mathbf{w}^\top\}_{u \in \mathcal{N}_{ST}(v, k, t)}$;
- 8 $\hat{\mathbf{A}}^{t(l)} \leftarrow \cup_{v \in \mathcal{V}^t} \{u \in \mathcal{N}_{ST}(v, k, t) \mid u \sim \text{Bern}(\phi_{v,k}^{t(l)})\}$
- 9 or $\cup_{v \in \mathcal{V}^t} \{u \in \mathcal{N}_{ST}(v, k, t) \mid u \sim \text{Cat}(\phi_{v,k}^{t(l)})\}$;
- 10 $\mathbf{Z}^{t(l)} \leftarrow \sum_{(u,v) \in \hat{\mathbf{A}}^{t(l)}} \{\hat{\mathbf{Z}}_v^{t(l-1)}\}_{v \in \mathcal{V}^t}$;
- 11 **end**
- 12 $\hat{\mathbf{Y}}^{T+1} = g(\mathbf{Z}^{T+1(L)})$;
- 13 $\text{DGIB}_{MS} \leftarrow \text{Eq. (8)}$, $\text{DGIB}_C \leftarrow \text{Eq. (13)}$;
- 14 Calculate the overall loss, as $\mathcal{L}_{\text{DGIB}} \leftarrow \text{Eq. (21)}$;
- 15 Update θ by minimizing $\mathcal{L}_{\text{DGIB}}$ and back-propagation.
- 16 **end**

Complexity Analysis. We analyze the computational complexity of each part in DGIB as follows. For brevity, denote $|\mathcal{V}|$ and $|\mathcal{E}|$ as the total number of nodes and edges in each graph snapshot, respectively. Let d be the dimension of the input node features, and d' be the dimension of the latent node representation.

- Linear input feature projection layer: $\mathcal{O}((T+1)|\mathcal{V}|dd')$.
- Relative Time Encoding (RTE(\cdot)) layer: $\mathcal{O}((T+1)|\mathcal{V}|d'^2)$.
- Spatio-temporal neighbor sampling and attention computation: $\mathcal{O}(L(T+1)|\mathcal{V}|) + \mathcal{O}(2Lk(T+1)|\mathcal{V}||\mathcal{E}|d')$, where k is the range of receptive field, and L is the number of layers.
- Feature aggregation: Constant complexity brought about by addition operations (ignored).
- Link prediction: $\mathcal{O}(|\mathcal{E}'|d')$, where \mathcal{E}' is the number of sampled links to be predicted.

In summary, the overall computational complexity of DGIB is:

$$\mathcal{O}(2Lk(T+1)|\mathcal{V}||\mathcal{E}|dd'^2) + \mathcal{O}(|\mathcal{E}'|d'). \quad (\text{A.1})$$

In conclusion, DGIB has a *linear* computation complexity with respect to the number of nodes and edges in all graph snapshots, which is on par with the state-of-the-art DGNNs. In addition, based on our experiments experience, DGIB can be trained and tested under the hardware configurations (including memory requirements) listed in Appendix D.2, and the training time consumption is listed in Appendix C.4, which demonstrates our DGIB has similar time complexity compared with most of the existing DGNNs.

B PROOFS

In this section, we provide proofs of Proposition 1, Proposition 2, and Proposition 3.

B.1 Proof of Proposition 1

PROOF. We apply the variational bounds of mutual information I_{NWJ} proposed by [28], which is thoroughly concluded in [31].

Lemma 1 (Mutual Information Variational Bounds in I_{NWJ}). Given any \mathbf{X} and \mathbf{Y} , and any permutation invariant function $f(\cdot, \cdot)$:

$$I(\mathbf{X}; \mathbf{Y}) \geq \mathbb{E}_{\mathbb{P}(\mathbf{X}, \mathbf{Y})} [f(\mathbf{X}, \mathbf{Y})] - e^{-1} \mathbb{E}_{\mathbb{P}(\mathbf{X})\mathbb{P}(\mathbf{Y})} [e^{f(\mathbf{X}, \mathbf{Y})}]. \quad (\text{B.1})$$

As $f(\mathbf{X}, \mathbf{Y})$ must learn to self-normalize, yielding a unique solution for variables $(\mathbf{Y}^{T+1}, \mathbf{Z}^{T+1})$ by plugging:

$$f(\mathbf{Y}^{T+1}, \mathbf{Z}^{T+1}) = 1 + \log \frac{\mathcal{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathcal{Q}_2(\mathbf{Y}^{T+1})} \quad (\text{B.2})$$

into Eq. (B.1). Specifically:

$$I(\mathbf{Y}^{T+1}; \mathbf{Z}^{T+1}) \geq \mathbb{E}_{\mathbb{P}(\mathbf{Y}^{T+1}, \mathbf{Z}^{T+1})} \left[1 + \log \frac{\mathcal{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathcal{Q}_2(\mathbf{Y}^{T+1})} \right] - e^{-1} \mathbb{E}_{\mathbb{P}(\mathbf{X})\mathbb{P}(\mathbf{Y})} \left[e \cdot \frac{\mathcal{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathcal{Q}_2(\mathbf{Y}^{T+1})} \right] \quad (\text{B.3})$$

$$\geq 1 + \mathbb{E}_{\mathbb{P}(\mathbf{Y}^{T+1}, \mathbf{Z}^{T+1})} \left[\log \frac{\mathcal{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathcal{Q}_2(\mathbf{Y}^{T+1})} \right] - \mathbb{E}_{\mathbb{P}(\mathbf{Y}^{T+1})\mathbb{P}(\mathbf{Z}^{T+1})} \left[\frac{\mathcal{Q}_1(\mathbf{Y}^{T+1} \mid \mathbf{Z}^{T+1})}{\mathcal{Q}_2(\mathbf{Y}^{T+1})} \right]. \quad (\text{B.4})$$

We conclude the proof of Proposition 1. \square

B.2 Proof of Proposition 2

PROOF. We apply the Data Processing Inequality (DPI) [3] and the Markovian dependency to prove the first inequality.

Lemma 2 (Mutual Information Lower Bound in Markov Chain). Given any three variables \mathbf{X} , \mathbf{Y} and \mathbf{Z} , which follow the Markov Chain $\langle \mathbf{X} \rightarrow \mathbf{Y} \rightarrow \mathbf{Z} \rangle$, we have:

$$I(\mathbf{X}; \mathbf{Y}) \geq I(\mathbf{X}; \mathbf{Z}). \quad (\text{B.5})$$

Directly applying Lemma 2 to the Markov Chain in DGIB, *i.e.*, $\langle \mathcal{D} \rightarrow (\{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z}) \rightarrow \mathbf{Z}^{T+1} \rangle$, which satisfies the Markov property $\mathcal{D} \perp\!\!\!\perp \mathbf{Z}^{T+1} \mid (\{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z})$, we have:

$$I(\mathcal{D}; \mathbf{Z}^{T+1}) \leq I(\mathcal{D}; \{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z}). \quad (\text{B.6})$$

Next, we prove the second inequality. To guarantee the compression order following ‘‘structure first, features second’’ in each time step, as well as the Spatio-Temporal Local Dependence (Assumption 2), we define the order:

Definition 2 (DGIB Markovian Decision Order). For any variables in set $\{\hat{\mathbf{A}}^t\}_{t \in \mathcal{I}_A} \cup \{\mathbf{Z}^t\}_{t \in \mathcal{I}_Z}$, the Markovian decision order \prec :

- For different time indices t_1 and t_2 , $\hat{\mathbf{A}}^{t_1}, \mathbf{Z}^{t_1} \prec \hat{\mathbf{A}}^{t_2}, \mathbf{Z}^{t_2}$.
- For any individual time index t , $\hat{\mathbf{A}}^t \prec \mathbf{Z}^t$.

To satisfy the Spatio-Temporal Local Dependence Assumption, we define two preceding sequences of sets based on the order $<$:

$$\mathcal{S}_A^t, \mathcal{S}_Z^t = \{\hat{A}^{t_1}, Z^{t_1} \mid t_1, t_2 < t - 1, t_1 \in \mathcal{I}_A, t_2 \in \mathcal{I}_Z\}. \quad (\text{B.7})$$

Thus, we have $\mathcal{S}_A^t \perp\!\!\!\perp \hat{A}^t \mid \{Z^{t-1}, A^t\}$ and $\mathcal{S}_Z^t \perp\!\!\!\perp Z^t \mid \{Z^{t-1}, \hat{A}^t\}$.

Then, we decompose $I(\mathcal{D}; \{\hat{A}^t\}_{t \in \mathcal{I}_A} \cup \{Z^t\}_{t \in \mathcal{I}_Z})$ into:

$$I(\mathcal{D}; \{\hat{A}^t\}_{t \in \mathcal{I}_A} \cup \{Z^t\}_{t \in \mathcal{I}_Z}) = \sum_{t \in \mathcal{I}_A} I(\mathcal{D}; \hat{A}^t \mid \mathcal{S}_A^t) + \sum_{t \in \mathcal{I}_Z} I(\mathcal{D}; Z^t \mid \mathcal{S}_Z^t). \quad (\text{B.8})$$

Following [46], we provide upper bounds for $I(\mathcal{D}; \hat{A}^t \mid \mathcal{S}_A^t)$ and $I(\mathcal{D}; Z^t \mid \mathcal{S}_Z^t)$, respectively:

$$\begin{aligned} I(\mathcal{D}; \hat{A}^t \mid \mathcal{S}_A^t) &\leq I(\mathcal{D}, Z^{t-1}; \hat{A}^t \mid \mathcal{S}_A^t) = I(Z^{t-1}, A^t; \hat{A}^t \mid \mathcal{S}_A^t) \\ &\leq I(Z^{t-1}, A^t; \hat{A}^t) = \mathcal{A}^t - \text{KL}[\mathbb{P}(\hat{A}^t) \parallel \mathbb{Q}(\hat{A}^t)] \\ &\leq \mathcal{A}^t. \end{aligned} \quad (\text{B.9})$$

Similarly, we have:

$$\begin{aligned} I(\mathcal{D}; Z^t \mid \mathcal{S}_Z^t) &\leq I(\mathcal{D}, Z^{t-1}, \hat{A}^t; Z^t \mid \mathcal{S}_Z^t) = I(Z^{t-1}, \hat{A}^t; Z^t \mid \mathcal{S}_Z^t) \\ &\leq I(Z^{t-1}, \hat{A}^t; Z^t) = \mathcal{Z}^t - \text{KL}[\mathbb{P}(Z^t) \parallel \mathbb{Q}(Z^t)] \\ &\leq \mathcal{Z}^t. \end{aligned} \quad (\text{B.10})$$

By plugging Eq. (B.9) and Eq. (B.10) into Eq. (B.8), we conclude the proof of the second inequality. \square

B.3 Proof of Proposition 3

PROOF. We apply the upper bound proposed in the Variational Information Bottleneck (VIB) [1].

Lemma 3 (Mutual Information Upper Bound in VIB). Given any two X and Y , we have the variational upper bound of $I(X; Y)$:

$$\begin{aligned} I(X; Y) &= \mathbb{E}_{\mathbb{P}(X, Y)} \left[\log \frac{\mathbb{P}(Y \mid X)}{\mathbb{P}(Y)} \right] = \mathbb{E}_{\mathbb{P}(X, Y)} \left[\log \frac{\mathbb{P}(Y \mid X) \mathbb{Q}(Y)}{\mathbb{P}(Y) \mathbb{Q}(Y)} \right] \\ &= \mathbb{E}_{\mathbb{P}(X, Y)} \left[\log \frac{\mathbb{P}(Y \mid X)}{\mathbb{Q}(Y)} \right] - \underbrace{\text{KL}[\mathbb{P}(Y) \parallel \mathbb{Q}(Y)]}_{\text{non-negative}} \\ &\leq \mathbb{E}_{\mathbb{P}(X, Y)} \left[\log \frac{\mathbb{P}(Y \mid X)}{\mathbb{Q}(Y)} \right]. \end{aligned} \quad (\text{B.11})$$

By utilizing Lemma 3 to variables $Z^{1:T}$ and Z^{T+1} , we conclude the proof of Proposition 3. \square

C EXPERIMENT DETAILS AND ADDITIONAL RESULTS

In this section, we provide additional experiment details and results.

C.1 Datasets Details

We use three real-world datasets to evaluate DGIB on the challenging future link prediction task.

- **COLLAB**¹ [41] is an academic collaboration dataset with papers that were published during 1990-2006 (16 graph snapshots), which reveals the dynamic citation networks among

¹<https://www.aminer.cn/collaboration>

authors. Nodes and edges represent authors and co-authorship, respectively. Based on the co-authored publication, there are five attributes in edges, including “Data Mining”, “Database”, “Medical Informatics”, “Theory” and “Visualization”. We apply word2vec [26] to extract 32-dimensional node features from paper abstracts.

- **Yelp**² [33] contains customer reviews on business, which are collected from the crowd-sourced local business review and social networking web. Nodes represent customer or business, and edges represent review behaviors, respectively. Considering categories of business, there are five attributes in edges, including “Pizza”, “American (New) Food”, “Coffee & Tea”, “Sushi Bars” and “Fast Food” from January 2019 to December 2020 (24 graph snapshots). We apply word2vec [26] to extract 32-dimensional node features from reviews.
- **ACT**³ [22] describes student actions on a popular MOOC website within a month (30 graph snapshots). Nodes represent students or targets of actions, edges represent actions. Considering the attributes of different actions, we apply K-Means to cluster the action features into five categories. We assign the features of actions to each student or target and expand the original 4-dimensional features to 32 dimensions by a linear function.

Statistics of the three datasets are concluded in Table B.1. These three datasets have different time spans and temporal granularity (16 years, 24 months, and 30 days), covering most real-world scenarios. The most challenging dataset for the future link prediction task is the COLLAB. In addition to having the longest time span and the coarsest temporal granularity, it also has the largest difference in the properties of its links, which greatly challenges the robustness.

Table B.1: Statistics of the real-world dynamic graph datasets.

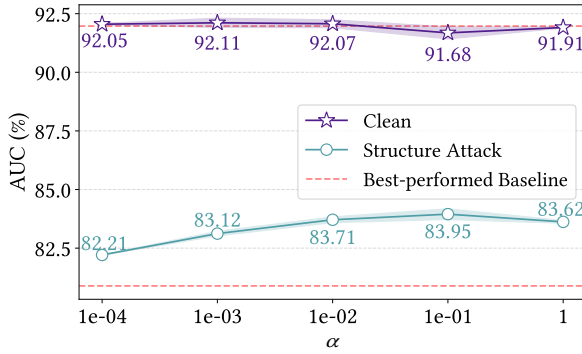
Dataset	# Node	# Link	# Link Type	Length (Split)	Temporal Granularity
COLLAB	23,035	151,790	5	16 (10/1/5)	year
Yelp	13,095	65,375	5	24 (15/1/8)	month
ACT	20,408	202,339	5	30 (20/2/8)	day

C.2 Detailed Settings and Analysis for Experiment in Figure 3

In Figure 3, we evaluate the performance of our DGIB against non-targeted adversarial attack compared with GIB [46]. We choose the ACT [22] as the dataset, with 20 graphs to train, 2 graphs to validate, and 8 graphs to test. In attacking settings, we follow the same non-targeted adversarial attack settings introduced in Section 5.1.3 on graph structures. For the baseline, we adapt the GIB [46]-Cat to dynamic scenarios by obtaining \hat{Z}^t for each individual graph snapshot first with the original GIB [46], and then aggregating them using the vanilla LSTM [15] to learn the comprehensive representation \hat{Z}^{T+1} , where the two steps are jointly optimized by the overall objective Eq. (3). For our DGIB, we train the DGIB-Cat version with the same objective in Eq. (3).

²<https://www.yelp.com/dataset>

³<https://snap.stanford.edu/data/act-mooc.html>

Figure B.1: Sensitivity analysis of α on COLLAB.

Analysis. Results demonstrates that directly optimizing GIB with the intuitive IB objective in Eq. (3) will lead to sub-optimal model performance, and our DGIB is endowed with stronger robustness by jointly optimizing $DGIB_{MS}$ and $DGIB_C$ channels, which cooperate and constrain each other to satisfy the *MSC* Condition.

C.3 Parameter Sensitivity Analysis

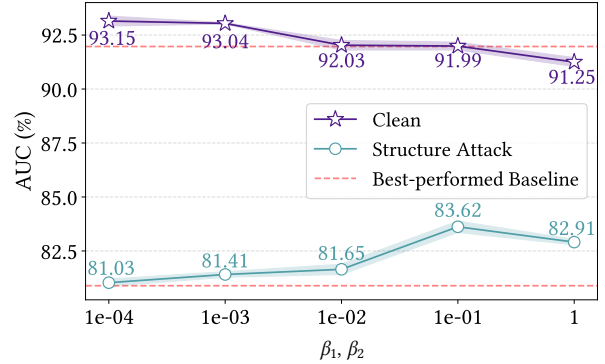
We provide additional experiments on the sensitivity of hyperparameters α , β_1 and β_2 , which are chosen from $\{1e-04, 1e-03, 1e-02, 1e-01, 1\}$. We report the results of sensitivity analysis on the clean and structure-attacked COLLAB in Figure B.1 and Figure B.2.

Results. Results demonstrate that the performance on both the clean and structure-attacked COLLAB is sensitive to different values of α , β_1 , and β_2 , and contains in a reasonable range. In addition, we observe there exist negative correlations between performance on the clean dataset and attacked dataset for three hyperparameters, which demonstrates the confrontations of the MSE term and compression term in DGIB overall optimization objectives (Eq. (6)). Specifically, in most cases, higher α , β_1 , and β_2 , lead to better robustness but weaker clean dataset performance. In conclusion, different combinations of hyperparameters lead to varying task performance and model robustness, and we follow the tradition of configuring the values of hyperparameters with the best trade-off setting against adversarial attacks.

C.4 Training Efficiency Analysis

We report the training time for our DGIB-Bern and DGIB-Cat with the default configurations as we provided in the code. We conduct experiments with the hardware and software configurations listed in Section D.2. We ignore the tiny difference when training under the poisoning attacks and only report the average time per epoch for training on the respective clean datasets in Table B.2.

Results. The training time of DGIB-Bern and DGIB-Cat is in the same order as the state-of-the-art DGNN baselines due to the computation complexity of DGIB is on par with related works. The maximum epochs for each training are set to 1000 (fixed), thus the total time is feasible in practice. Note that, the number of layers, neighbors to sampling, *etc.*, will have a significant impact on the training efficiency, and not always large numbers bring extra

Figure B.2: Sensitivity analysis of β_1, β_2 on COLLAB.

improvements in performance, so it is recommended to properly set the related parameters.

Table B.2: Results of training time per epoch (s).

Dataset	COLLAB	Yelp	ACT
DIDA [56]	2.61	5.12	3.31
DGIB-Bern	0.88	1.91	1.32
DGIB-Cat	0.86	1.54	1.64

D IMPLEMENTATION DETAILS

In this section, we provide implementation details of the DGIB.

D.1 DGIB Implementation Details

According to respective settings, we randomly split the dynamic graph datasets into training, validation, and testing chronological sets. We sample negative links from nodes that do not have links, and the negative links for validation and testing are kept the same for all baseline methods and ours. We set the number of positive links to the same as the negative links. We use the Area under the ROC Curve (AUC) as the evaluation metric. As we focus on the future link prediction task, we use the inner product of a pair of learned node representations to predict the occurrence of links, *i.e.*, we implement the link predictor $g(\cdot)$ as the inner product of hidden embeddings, which is widely applied in future link prediction tasks. The non-linear rectifier τ is $\text{ReLU}(\cdot)$ [27], and the activation function σ is $\text{Sigmoid}(\cdot)$. We randomly run all the experiments five times, and report the average results with standard deviations.

D.2 Hardware and Software Configurations

We conduct the experiments with:

- Operating System: Ubuntu 20.04 LTS.
- CPU: Intel(R) Xeon(R) Platinum 8358 CPU@2.60GHz with 1TB DDR4 of Memory.
- GPU: NVIDIA Tesla A100 SMX4 with 40GB of Memory.
- Software: CUDA 10.1, Python 3.8.12, PyTorch⁴ 1.9.1, PyTorch Geometric⁵ 2.0.1.

More discussions in the *OpenReview* forum.

⁴<https://github.com/pytorch/pytorch>

⁵https://github.com/pyg-team/pytorch_geometric