
Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts

Zeyang Zhang^{1*}, Xin Wang^{1†}, Ziwei Zhang¹, Zhou Qin²,
Weigao Wen², Hui Xue², Haoyang Li¹, Wenwu Zhu^{1†}

¹Department of Computer Science and Technology, BNRist, Tsinghua University, ²Alibaba Group
zy-zhang20@mails.tsinghua.edu.cn, {xin_wang, zwzhang}@tsinghua.edu.cn,
{qinzhou.qinzhou, weigao.wen, hui.xue}@alibaba-inc.com,
lihy18@mails.tsinghua.edu.cn, wwzhu@tsinghua.edu.cn

Abstract

Dynamic graph neural networks (DyGNNs) currently struggle with handling distribution shifts that are inherent in dynamic graphs. Existing work on DyGNNs with out-of-distribution settings only focuses on the time domain, failing to handle cases involving distribution shifts in the spectral domain. In this paper, we discover that there exist cases with distribution shifts unobservable in the time domain while observable in the spectral domain, and propose to study distribution shifts on dynamic graphs in the spectral domain for the first time. However, this investigation poses two key challenges: i) it is non-trivial to capture different graph patterns that are driven by various frequency components entangled in the spectral domain; and ii) it remains unclear how to handle distribution shifts with the discovered spectral patterns. To address these challenges, we propose Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts (**SILD**), which can handle distribution shifts on dynamic graphs by capturing and utilizing invariant and variant spectral patterns. Specifically, we first design a DyGNN with Fourier transform to obtain the ego-graph trajectory spectrums, allowing the mixed dynamic graph patterns to be transformed into separate frequency components. We then develop a disentangled spectrum mask to filter graph dynamics from various frequency components and discover the invariant and variant spectral patterns. Finally, we propose invariant spectral filtering, which encourages the model to rely on invariant patterns for generalization under distribution shifts. Experimental results on synthetic and real-world dynamic graph datasets demonstrate the superiority of our method for both node classification and link prediction tasks under distribution shifts.

1 Introduction

Dynamic graph neural networks (DyGNNs) have achieved remarkable success in many predictive tasks over dynamic graphs [1, 2]. Existing DyGNNs exhibit limitations in handling distribution shifts, which naturally exist in dynamic graphs due to multiple uncontrollable factors [3, 4, 5, 6]. Existing work on out-of-distribution generalized DyGNNs focuses on handling distribution shifts in the time domain. For example, DIDA [7] utilizes dynamic graph attention to mask the graph trajectories to capture the invariant patterns on dynamic graphs, which assumes that in the time domain, the distribution shift is observable and the invariant and variant patterns can be easily disentangled.

However, there exist cases that the distribution shift is unobservable in the time domain while observable in the spectral domain, as shown in Figure 1. The shift in frequency components can

*This work was done during the author’s internship at Alibaba Group

†Corresponding authors




Figure 1: An illustration example: the graph dynamics from different frequency components are entangled in the temporal domain, while it is much easier to distinguish different frequency components by masking the spectrums in the spectral domain. In this case, the frequency components in the invariant spectrums determine the node labels, while the relationship between the variant spectrums and labels is not stable under distribution shifts.

be clearly observed in the spectral domain, while these components are indistinguishable in the time domain. Moreover, in real-world applications, the observed dynamic graphs usually consist of multiple mixed graph structural and featural dynamics from various frequency components [8, 9, 10].

To address this problem, in this paper, we study the problem of handling distribution shifts on dynamic graphs in the spectral domain for the first time, which poses the following two key challenges: i) it is non-trivial to capture different graph patterns that are driven by various frequency components entangled in the spectral domain, and ii) it remains unclear how to handle distribution shifts with the discovered spectral patterns.

To tackle these challenges, we propose Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts (**SILD**³). Our proposed **SILD** model can effectively handle distribution shifts on dynamic graphs by discovering and utilizing invariant and variant spectral patterns. Specifically, we first design a DyGNN with Fourier transform to obtain the ego-graph trajectory spectrums so that the mixed graph dynamics can be transformed into separate frequency components. Then we develop a disentangled spectrum mask that leverages the amplitude and phase information of the ego-graph trajectory spectrums to obtain invariant and variant spectrum masks so that graph dynamics from various frequency components can be filtered. Finally, we propose an invariant spectral filtering that discovers the invariant and variant patterns via the disentangled spectrum masks, and minimize the variance of predictions with exposure to various variant patterns. As such, **SILD** is able to exploit invariant patterns to make predictions under distribution shifts. Experimental results on several synthetic and real-world datasets, including both node classification and link prediction tasks, demonstrate the superior performance of our **SILD** model compared to state-of-the-art baselines under distribution shifts. To summarize, we make the following contributions:

- We propose to study distribution shifts on dynamic graphs in the spectral domain, to the best of our knowledge, for the first time.
- We propose Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts (**SILD**), which can handle distribution shifts on dynamic graphs in the spectral domain.
- We employ DyGNN with Fourier transform to obtain the node spectrums, design a disentangled spectrum mask to obtain invariant and variant spectrum masks in the spectral domain, and propose the invariant spectral filtering mechanism so that **SILD** is able to handle distribution shifts.
- We conduct extensive experiments on several synthetic and real-world datasets, including both node classification and link prediction tasks, to demonstrate the superior performance of our method compared to state-of-the-art baselines under distribution shifts.

³The codes are available at Github.

2 Problem Formulation and Notations

Dynamic Graphs A dynamic graph can be represented as $\mathcal{G} = (\{\mathcal{G}^t\}_{t=1}^T)$, where T represents the total number of time stamps, and each $\mathcal{G}^t = (\mathcal{V}^t, \mathcal{E}^t)$ corresponds to a graph snapshot at time stamp t with the node set \mathcal{V}^t and the edge set \mathcal{E}^t . For simplicity, we also represent a graph snapshot as $\mathcal{G}^t = (\mathbf{X}^t, \mathbf{A}^t)$, which includes the node feature matrix \mathbf{X}^t and the adjacency matrix \mathbf{A}^t . We further denote a random variable of \mathcal{G}^t as \mathbf{G}^t . The prediction task on dynamic graphs aims to utilize past graph snapshots to make predictions, *i.e.*, $p(\mathbf{Y}^t | \mathbf{G}^{1:t})$, where $\mathbf{G}^{1:t} = \{\mathbf{G}^1, \mathbf{G}^2, \dots, \mathbf{G}^t\}$ denotes the graph trajectory, and the label \mathbf{Y}^t represent the node properties or the links at time $t + 1$. For brevity, we take node-level prediction tasks as an example in this paper. Following [7], the distribution of graph trajectory can be factorized into ego-graph trajectories, such that $p(\mathbf{Y}^t | \mathbf{G}^{1:t}) = \prod_v p(\mathbf{y}_v^t | \mathbf{G}_v^{1:t})$.

Distribution Shifts on Dynamic Graphs The common optimization objective for prediction tasks on dynamic graphs is to learn an optimal predictor with empirical risk minimization (ERM), *i.e.* $\min_{\theta} \mathbb{E}_{(y^t, \mathcal{G}_v^{1:t}) \sim p_{tr}(\mathbf{y}^t, \mathbf{G}_v^{1:t})} \mathcal{L}(f_{\theta}(\mathcal{G}_v^{1:t}), y^t)$, where f_{θ} is a learnable dynamic graph neural networks. Under distribution shifts, however, the optimal predictor trained with ERM and the training distribution may not generalize well to the test distribution, since the risk minimization objectives under two distributions are different due to $p_{tr}(\mathbf{Y}^t, \mathbf{G}^{1:t}) \neq p_{te}(\mathbf{Y}^t, \mathbf{G}^{1:t})$. The distribution shift on dynamic graphs is complex that may originate from temporal distribution shifts [11, 6, 12, 13, 14] as well as structural distribution shifts [15, 16, 17]. For example, trends or community structures can affect interaction patterns in co-author networks [18] and recommendation networks [19], *i.e.*, the distribution of ego-graph trajectories may vary through time and structures.

Following out-of-distribution (OOD) generalization literature [7, 15, 11, 20, 21, 22], we make the following assumptions of distribution shifts on dynamic graphs:

Assumption 1 For a given task, there exists a predictor $f(\cdot)$, for samples $(\mathcal{G}_v^{1:t}, y^t)$ from any distribution, there exists an invariant pattern $P_I^t(v)$ and a variant pattern $P_V^t(v)$ such that the following conditions are satisfied: 1) the invariant patterns are sufficient to predict the labels, $y_v^t = f(P_I^t(v)) + \epsilon$, where ϵ is a random noise, 2) the observed data is composed of invariant and variant patterns, $P_I^t(v) = \mathcal{G}_v^{1:t} \setminus P_V^t(v)$, 3) the influence of the variant patterns on labels is shielded by the invariant patterns, $\mathbf{y}_v^t \perp \mathbf{P}_V^t(v) | \mathbf{P}_I^t(v)$.

In the next section, inspired by [16], we give a motivation example to provide some high-level intuition before going to our formal method.

3 Motivation Example

Here we introduce a toy dynamic graph example to motivate learning invariant patterns in the spectral domain. We assume that the invariant and variant patterns lie in the 1-hop neighbors, *i.e.*, each node has an invariant subgraph and a variant subgraph. For simplicity, we focus on the number of neighbors, *i.e.*, each node v has an invariant subgraph related degree $\mathbf{d}_{v,1} \in \mathbb{R}^{T \times 1}$ and a variant subgraph related degree $\mathbf{d}_{v,2} \in \mathbb{R}^{T \times 1}$. Only the former determines the node label, *i.e.*, $y_v = \mathbf{g}^{\top} \mathbf{d}_{v,1}$. Note that invariant and variant subgraphs are not observed in the data. We further assume a one-dimensional constant feature for each node, which is set as 1 without loss of generality.

For the model in the spatial-temporal domain, we adopt sum pooling as one-layer graph convolution, *i.e.*, the message passing for each node and time is $\mathbf{h}_v = \sum_{u \in \mathcal{N}_v} 1 = \mathbf{d}_{v,1} + \mathbf{d}_{v,2}$. We further adopt a mask $\mathbf{m} \in \mathbb{R}^{T \times 1}$ to filter patterns in the temporal domain and make predictions by a linear classifier, *i.e.*, $\hat{y}_v = \mathbf{w}^{\top} (\mathbf{m} \odot \mathbf{h}_v)$, where $\mathbf{w} \in \mathbb{R}^{T \times 1}$ denotes the learnable parameters. Then, the empirical risk in the training dataset D_{tr} is $R_{tr}(\mathbf{w}) = \frac{1}{|D_{tr}|} \sum_{v \in D_{tr}} (\hat{y}_v - y_v)^2$. We have the following proposition.

Proposition 1 For any mask $\mathbf{m} \in \mathbb{R}^{T \times 1}$, for the optimal classifier in the training data $\mathbf{w}^* = \arg \min_{\mathbf{w}} R_{tr}(\mathbf{w})$, if $\|\mathbf{m} \odot \mathbf{w}^*\|_2 \neq 0$, there exist OOD nodes with unbounded error, *i.e.*, $\exists v$ s.t. $\lim_{\|\mathbf{d}_{v,2}\| \rightarrow \infty} (\hat{y}_v - y_v)^2 = \infty$.

The proposition 1 shows that a classifier trained with masks and empirical risk minimization has unbounded risks in testing data under distribution shifts as the classifier uses variant patterns to make predictions. Next, we show that under mild conditions, an invariant linear classifier in the spectral

domain can solve this problem. Denote $\Phi \in \mathbb{C}^{T \times T}$ as the Fourier bases, where $\Phi_{k,t} = \frac{1}{\sqrt{T}} e^{-j \frac{2\pi kt}{T}}$. Denote $\mathbf{z}_v = \Phi \sum_{u \in \mathcal{N}_v} \mathbf{1}$ as the spectral representation after a linear message-passing. The prediction is $\hat{y}_v = \mathbf{w}^H (\mathbf{m} \odot \mathbf{z}_v)$, where $\mathbf{m} \in \mathbb{C}^{T \times 1}$ is the mask to filter the spectral patterns, $\mathbf{w} \in \mathbb{C}^{T \times 1}$ is a linear classifier, and $(\cdot)^H$ denotes Hermitian transpose. We have the following proposition.

Proposition 2 *If $(\overline{\Phi \mathbf{d}_{v,1}} \odot \Phi \mathbf{d}_{v,1}) \odot (\overline{\Phi \mathbf{d}_{v,2}} \odot \Phi \mathbf{d}_{v,2}) = \mathbf{0}, \forall \mathbf{d}_{v,1}, \mathbf{d}_{v,2}$, then $\exists \mathbf{m} \in \mathbb{C}^{T \times 1}$ such that the optimal spectral classifier in the training data has bounded error; i.e., for $\mathbf{w}^* = \arg \min_{\mathbf{w}} R_{tr}(\mathbf{w})$, $\exists \epsilon > 0, \forall v, \lim_{\|\mathbf{d}_{v,2}\| \rightarrow \infty} (\hat{y}_v - y_v)^2 < \epsilon$.*

The proposition 2 shows that if the frequency bandwidths of invariant and variant patterns do not have any overlap, there exists a spectral mask such that a linear classifier trained with empirical risk minimization in the spectral domain will have bounded risk in any testing data distribution. This example motivates us to capture invariant and variant patterns in the spectral domain, which is not feasible in the spatial-temporal domain.

4 Method

In this section, we introduce our method named Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts (**SILD**) to handle distribution shifts in dynamic graphs, including three modules, dynamic graph neural networks with Fourier transform, disentangled spectrum mask, and invariant spectral filtering. The framework of our method is shown in Figure 2.

4.1 Dynamic Graph Neural Network with Spectral Transform

Dynamic Graph Trajectories Modeling Each node on the dynamic graph has its ego-graph trajectory evolving through time that may determine the node properties or the occurrence of future links. Following [23, 24, 25], we adopt a message-passing network for each graph snapshot to aggregate the neighborhood information at the current time, i.e.,

$$\mathbf{m}_{u \rightarrow v}^t \leftarrow \text{MSG}(\mathbf{h}_u^t, \mathbf{h}_v^t), \mathbf{h}_v^t \leftarrow \text{AGG}(\{\mathbf{m}_{u \rightarrow v}^t \mid u \in \mathcal{N}^t(v)\}, \mathbf{h}_v^t), \quad (1)$$

where ‘MSG’ and ‘AGG’ denote message and aggregation functions, $\mathbf{m}_{u \rightarrow v}^t$ is the message from node u to node v , \mathbf{h}_u^t is the node embedding for node u at time t , $\mathcal{N}^t(v) = \{u \mid (u, v) \in \mathcal{E}^t\}$ is node v ’s neighborhood at time t . To model the high-order neighborhood information, we can stack multiple message-passing layers. In this way, the node embedding along time $\{\mathbf{h}_u^t\}_{t=1}^T$ summarizes the evolution of node u ’s ego-graph trajectories. We denote $\mathbf{H} \in \mathbb{R}^{T \times N \times d}$ as the ego-graph trajectory signals for all nodes on the dynamic graph, where T denotes the total time length, N denotes the number of nodes and d denotes the hidden dimensionality.

Spectral Transform As some patterns on dynamic graphs are unobservable in the time domain, while observable in the spectral domain, we transform the summarized ego-graph trajectory signals \mathbf{H} into the spectral domain via Fourier transform for each node and hidden dimension, i.e.,

$$\Phi_{k,t} = \frac{1}{\sqrt{T}} e^{-j \frac{2\pi kt}{T}}, \mathbf{Z} = \Phi \mathbf{H}, \quad (2)$$

where $\Phi \in \mathbb{C}^{K \times T}$ denotes the Fourier bases, K denotes the number of frequency components, and $\mathbf{Z} \in \mathbb{C}^{K \times N \times d}$ denote the node embeddings along frequency components in the spectral domain, and $\mathbf{Z}_{k,n,m} = \sum_{t=1}^T \Phi_{k,t} \mathbf{H}_{t,n,m}$. By choosing the Fourier bases, our spectral transform has the following advantages: 1) we can use fast Fourier transform (FFT) [26] to accelerate the computation. The computation complexity of Eq. (2) can be reduced from $O(NdT^2)$ to $O(NdT \log T)$. 2) Each basis has clear semantics, e.g., \mathbf{Z}_k denotes the node embeddings at the k -th frequency component in the spectral domain. In this way, we can observe how the nodes on the dynamic graph evolve in different frequency bandwidths. 3) Fourier transform is able to capture global and periodic patterns [27, 28], which are common in real-world dynamic graphs, e.g., the interactions on e-commerce networks may result from seasonal sales or product service cycles.

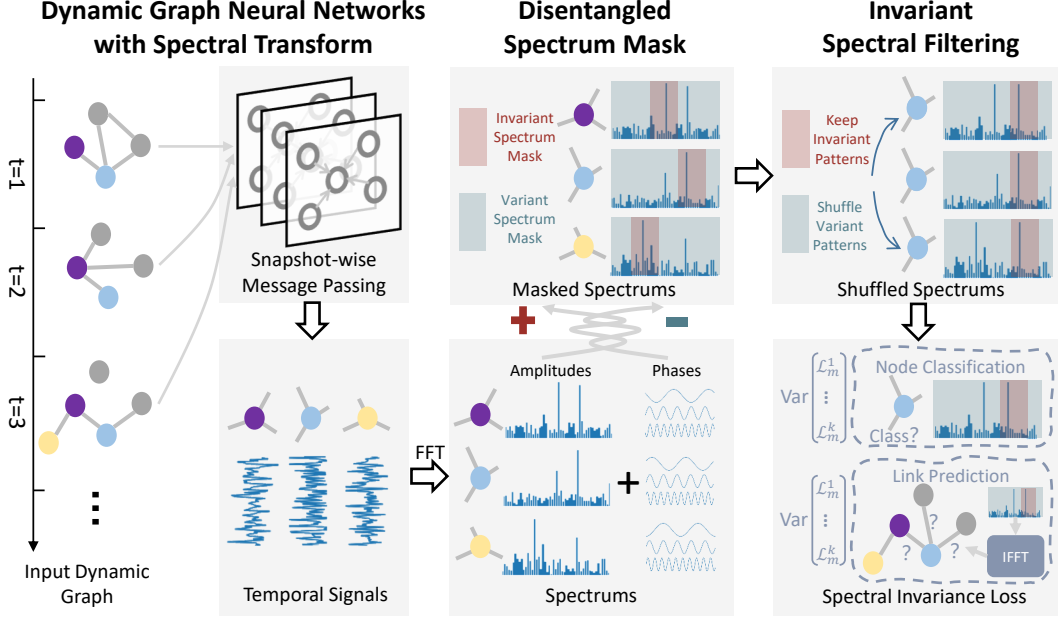


Figure 2: The framework of our proposed method **SILD**. Given a dynamic graph evolving through time, the dynamic graph neural networks with spectral transform first obtain the ego-graph trajectory spectrums in the spectral domain. Then the disentangled spectrum mask leverages the amplitude and phase information of the ego-graph trajectory spectrums to obtain invariant and variant spectrum masks. Last, invariant spectral filtering discovers the invariant and variant patterns via the disentangled spectrum masks, and minimizes the variance of predictions with exposure to various variant patterns, to help the model exploit invariant patterns to make predictions under distribution shifts.

4.2 Disentangled Spectrum Mask

To capture invariant patterns in the spectral domain, we propose to explicitly learn spectrum masks to disentangle the invariant and variant patterns. The embeddings in the spectral domain contain the amplitude information as well as the phase information for each node

$$\text{Amp}(\mathbf{Z}) = \left(\text{Imag}^2(\mathbf{Z}) + \text{Real}^2(\mathbf{Z}) \right)^{\frac{1}{2}}, \phi(\mathbf{Z}) = \arctan \frac{\text{Imag}(\mathbf{Z})}{\text{Real}(\mathbf{Z})}, \quad (3)$$

where $\text{Real}(\cdot)$ and $\text{Imag}(\cdot)$ denote the real and imaginary part of the complex number, *i.e.*, $\mathbf{Z} = \text{Real}(\mathbf{Z}) + j\text{Imag}(\mathbf{Z})$, j denotes the imaginary unit, $\text{Amp}(\mathbf{Z}) \in \mathbb{R}^{K \times N \times d}$ and $\phi(\mathbf{Z}) \in \mathbb{R}^{K \times N \times d}$ denote the amplitude and phase information respectively. For brevity, the tensor operators in Eq. (3) are all element-wise, *e.g.*, $(\text{Imag}^2(\mathbf{Z}))_{i,j,k} = (\text{Imag}(\mathbf{Z}_{i,j,k}))^2$. Then, we obtain the spectrum masks by leveraging both the amplitude and phase information

$$\mathbf{M} = \text{MLP}(\text{Real}(\mathbf{Z}) \parallel \text{Imag}(\mathbf{Z})), \mathbf{M}_I = \text{sigmoid}(\mathbf{M}/\tau), \mathbf{M}_V = \text{sigmoid}(-\mathbf{M}/\tau), \quad (4)$$

where MLP denotes multi-layer perceptrons, τ is the temperature, $\mathbf{M}_I \in [0, 1]^K$ and $\mathbf{M}_V \in [0, 1]^K$ denote the spectrum mask for invariant and variant patterns, and \parallel represents the concatenation of the embeddings. In this way, the invariant and variant masks have a negative relationship, and each node can have its own spectrum mask. As the phase information includes high-level semantics in the original signals [29, 30, 31, 32, 33], we keep the phase information unchanged to reduce harm in the fine-grained semantic information for the graph trajectories, and filter the spectrums by the learned disentangled masks in terms of amplitudes,

$$\mathbf{Z}_I = \left(\mathbf{M}_I \odot \text{Amp}(\mathbf{Z}) \right) \odot (\cos \phi(\mathbf{Z}) + j \sin \phi(\mathbf{Z})), \mathbf{Z}_V = \left(\mathbf{M}_V \odot \text{Amp}(\mathbf{Z}) \right) (\cos \phi(\mathbf{Z}) + j \sin \phi(\mathbf{Z})), \quad (5)$$

where \mathbf{Z}_I and \mathbf{Z}_V denote the summarized invariant and variant patterns in the spectral domain. For node classification tasks, we can directly adopt the spectrums for the classifier to predict classes. For link prediction tasks, we can utilize inverse fast Fourier transform (IFFT) to transform the embeddings into the temporal domain for future link prediction

$$\mathbf{H}'_I = \Phi^H \mathbf{Z}_I, \mathbf{H}'_V = \Phi^H \mathbf{Z}_V, \quad (6)$$

where $(\cdot)^H$ is Hermitian transpose, \mathbf{H}'_I and \mathbf{H}'_V denote the filtered invariant and variant patterns that are transformed back into the temporal domain respectively.

4.3 Invariant Spectral Filtering

Under distribution shifts, the variant patterns on dynamic graphs have varying relationships with labels, while the invariant patterns have sufficient predictive abilities with regard to labels. We propose invariant spectral filtering to capture the invariant and variant patterns in the spectral domain, and help the model focus on invariant patterns to make predictions, thus handling distribution shifts. We take node classification tasks for an example as follows.

Let $\mathbf{Z}_I \in \mathbb{C}^{K \times N \times d}$ and $\mathbf{Z}_V \in \mathbb{C}^{K \times N \times d}$ be the filtered invariant and variant spectrums in the spectral domain. Then we can utilize the invariant and variant node spectrums to calculate the task loss

$$\mathcal{L}_I = l(f_I(\mathbf{Z}_I), \mathbf{Y}), \mathcal{L}_V = l(f_V(\mathbf{Z}_V), \mathbf{Y}), \quad (7)$$

where $f_I(\cdot)$ and $f_V(\cdot)$ are the classifiers for invariant and variant patterns respectively, \mathbf{Y} is the labels, and l is the loss function. The task loss is utilized to capture the patterns with the predictive abilities of labels. Recall in Assumption 1, the influence of variant patterns on labels is shielded given invariant patterns as the invariant patterns have sufficient predictive abilities w.r.t labels, and thus the model’s predictions should not change when being exposed to different variant patterns and the original invariant patterns. Inspired by [15, 7], we calculate the invariance loss by

$$\mathcal{L}_{INV} = \text{Var}(\{\mathcal{L}_m \mid \tilde{\mathbf{z}} : \tilde{\mathbf{z}} \in \mathcal{S}\}), \quad (8)$$

where $\mathcal{L}_m \mid \tilde{\mathbf{z}}$ denotes the mixed loss to measure the model’s prediction ability with exposure to the specific variant pattern $\tilde{\mathbf{z}} \in \mathbb{C}^{K \times d}$ that is sampled from a set of variant patterns \mathcal{S} . We adopt all the node embeddings in \mathbf{Z}_V to construct the set of variant patterns \mathcal{S} . Inspired by [34, 15], we calculate the mixed loss as

$$\mathcal{L}_m \mid \tilde{\mathbf{z}} = l(f_I(\mathbf{Z}_I) \odot \sigma(f_V(\tilde{\mathbf{z}})), \mathbf{Y}), \quad (9)$$

where σ denotes the sigmoid function. Then, the final training objective is

$$\min_{\theta, f_I} \mathcal{L}_I + \lambda \mathcal{L}_{INV} + \min_{f_V} \mathcal{L}_V, \quad (10)$$

where θ is the parameters that encompass all the model parameters except the classifiers, λ is a hyperparameter to balance the trade-off between the model’s predictive ability and invariance properties. A larger λ encourages the model to capture patterns with better invariance under distribution shifts, with the potential risk of lower predictive ability during training, as the shortcuts brought by the variant patterns might be discarded in the training process. After training, we only adopt invariant patterns to make predictions in the inference stage. The overall algorithm for training on node classification datasets is summarized in Algo. 1.

Algorithm 1 Training pipeline for SILD on node classification datasets

Require: Training epochs L , sample number S , hyperparameter λ

- 1: **for** $l = 1, \dots, L$ **do**
 - 2: Obtain the node embeddings \mathbf{H} with snapshot-wise message passing as Eq. (1)
 - 3: Transform the node embeddings into the spectral domain with FFT as Eq. (2)
 - 4: Calculate the disentangled spectrum masks as Eq. (4)
 - 5: Filter spectrums into invariant and variant patterns as Eq. (5)
 - 6: Calculate the task loss as Eq. (7)
 - 7: Sample S variant patterns from collections of \mathbf{Z}_V and calculate the invariance loss as Eq. (8)
 - 8: Update the model according to Eq. (10)
 - 9: **end for**
-

5 Experiments

In this section, we conduct extensive experiments to verify that our proposed method can handle distribution shifts on dynamic graphs by discovering and utilizing invariant patterns in the spectral domain. More details of the settings and other results can be found in Appendix.

Baselines. We adopt several representative dynamic GNNs and Out-of-Distribution(OOD) generalization methods as our baselines:

Table 1: Results of different methods on real-world link prediction and node classification datasets. The best results are in bold and the second-best results are underlined. The year in the Aminer dataset denotes the test split, e.g., ‘Aminer15’ denotes the average test accuracy in 2015.

Task Dataset	Link Prediction (AUC%)		Node Classification (ACC%)		
	Collab	Yelp	Aminer15	Aminer16	Aminer17
GCRN	69.72±0.45	54.68±7.59	47.96±1.12	51.33±0.62	42.93±0.71
EGCN	76.15±0.91	53.82±2.06	44.14±1.12	46.28±1.84	37.71±1.84
DySAT	76.59±0.20	66.09±1.42	48.41±0.81	49.76±0.96	42.39±0.62
IRM	75.42±0.87	56.02±16.08	48.44±0.13	50.18±0.73	42.40±0.27
VREx	76.24±0.77	66.41±1.87	48.70±0.73	49.24±0.27	42.59±0.37
GroupDRO	76.33±0.29	66.97±0.61	48.73±0.61	49.74±0.26	42.80±0.36
DIDA	81.87±0.40	75.92±0.90	50.34±0.81	51.43±0.27	44.69±0.06
SILD	84.09±0.16	78.65±2.22	52.35±1.04	54.11±0.62	45.54±1.19

Table 2: Results of different methods on synthetic link prediction and node classification datasets. The best results are in bold and the second-best results are underlined. A larger ‘shift’ denotes a higher distribution shift level.

Dataset Shift	Link-Synthetic (AUC%)			Node-Synthetic (ACC%)		
	0.4	0.6	0.8	0.4	0.6	0.8
GCRN	72.57±0.72	72.29±0.47	67.26±0.22	27.19±2.18	25.95±0.80	29.26±0.69
EGCN	69.00±0.53	62.70±1.14	60.13±0.89	24.01±2.29	22.75±0.96	24.98±1.32
DySAT	70.24±1.26	64.01±0.19	62.19±0.39	40.95±2.89	37.94±1.01	30.90±1.97
IRM	69.40±0.09	63.97±0.37	62.66±0.33	33.23±4.70	30.29±1.71	29.43±1.38
VREx	70.44±1.08	63.99±0.21	62.21±0.40	41.78±1.30	38.11±2.81	29.56±0.44
GroupDRO	70.30±1.23	64.05±0.21	62.13±0.35	41.35±2.19	35.74±3.93	31.03±1.24
DIDA	85.20±0.84	82.89±0.23	72.59±3.31	43.33±7.74	39.48±7.93	28.14±3.07
SILD	85.95±0.18	84.69±1.18	78.01±0.71	43.62±2.74	39.78±3.56	38.64±2.76

- Dynamic GNNs: **GCRN** [35] is a representative dynamic GNN that first adopts a GCN[36] to obtain node embeddings and then a GRU [37] to model the network evolution. **EGCN** [25] adopts an LSTM [38] or GRU [37] to flexibly evolve the GCN [36] parameters through time. **DySAT** [24] aggregates neighborhood information at each graph snapshot using structural attention and models network dynamics with temporal self-attention.
- OOD generalization methods: **IRM** [20] aims at learning an invariant predictor which minimizes the empirical risks for all training domains. **GroupDRO** [39] puts more weight on training domains with larger errors to minimize the worst-group risks across training domains. **V-REx** [40] reduces the differences in the risks across training domains to reduce the model’s sensitivity to distributional shifts. As these methods are not specifically designed for dynamic graphs, we adopt the best-performed dynamic GNNs as their backbones on each dataset.
- OOD generalization methods for dynamic graphs: **DIDA** [7] utilizes disentangled attention to capture invariant and variant patterns in the spatial-temporal domain, and conducts spatial-temporal intervention mechanism to let the model focus on invariant patterns to make predictions.

5.1 Real-world Datasets

Settings We use 3 real-world dynamic graph datasets, including Collab [41, 7], Yelp [24, 7] and Aminer [42, 43]. Following [7], we adopt the challenging inductive future link prediction task on Collab and Yelp, where the model should exploit historical graphs to predict the occurrence of links in the next time step. To measure the model’s performance under distribution shifts, the model is tested on another dynamic graph with different fields, which is unseen during training. For node classification, we adopt Aminer, a citation network, where nodes represent papers, and edges from u to v with timestamp t denote the paper u published at year t cites the paper v . The task is to predict the venues of the papers. We train on papers published between 2001 - 2011, validate on

those published in 2012 - 2014, and test on those published since 2015. On this dataset, the model is tested to exploit the invariant patterns and make stable predictions under distribution shifts, where the patterns on the dynamic graph may vary in different years.

Results Based on the results in Table 1, we have the following observations: 1) *Under distribution shifts, the general OOD generalization baselines have limited improvements over the dynamic GNNs, e.g., GroupDRO improves over DySAT with 0.9% in Yelp and 0.3% in Aminer15 respectively. A plausible reason is that they are not specially designed to handle distribution shifts on dynamic graphs, and may not consider the graph structural and temporal dynamics to capture invariant patterns. Another reason might be that they strongly rely on high-quality environment labels to capture invariant patterns, which are almost unavailable on real-world dynamic graphs.* 2) *Our method can better handle distribution shifts than the baselines.* The datasets have strong distribution shifts, e.g., COVID-19 happens midway and has considerable influence on the consumer behavior on Yelp, and the citation patterns may shift with the outbreak of deep neural networks on Aminer. Nevertheless, our method **SILD** has significant improvements over the state-of-the-art OOD generalization baseline for dynamic graphs DIDA on all datasets, e.g., 2% on average for most datasets, which verifies that our method can better capture the invariant and variant patterns in the spectral domain, and thus handling distribution shifts on dynamic graphs.

5.2 Synthetic Datasets

Settings To evaluate the model’s generalization ability under distribution shifts, we conduct experiments on synthetic link prediction and node classification datasets, which are constructed by introducing manually-designed distribution shifts. For link prediction datasets, we follow [7] to generate additional varying features for each node and timestamps on the original dataset Collab, where these additional features are constructed with spurious correlations w.r.t the labels, *i.e.*, the links in the next timestamps. The spurious correlation degree is determined by a shift level parameter. On this dataset, to have better generalization ability, the model should not rely on variant patterns that exploit the additional features with spurious correlations. For node classification, we briefly introduce the construction of the synthetic dataset as follows. We generate the dynamic graph with stochastic block model [44], where the link probability between nodes at each graph snapshot is determined by two frequency factors. The correlation of one of the factors with class labels is always 1, while the other factor has a variant relationship with labels, where the relationship is also controlled by a shift level parameter. The model should discover and focus on the invariant frequency factors whose relationship with labels is invariant under distribution shifts. For both datasets, we set the shift level parameters as 0.4, 0.6, 0.8 for training and validation splits, and 0 for test splits.

Results Based on the results in Table 2, we have the following observations: 1) *Our method can better handle distribution shifts than the baselines, especially under stronger distribution shifts.* **SILD** consistently outperforms DyGNN and general OOD generalization baselines by a significantly large margin, which can credit to our special design to handle distribution shifts on dynamic graphs in the spectral domain. Our method also has a significant improvement over the best-performed baseline under the strongest distribution shift, e.g., with absolute improvements of 5% in Link-Synthetic(0.8) and 7% in Node-Synthetic(0.8) respectively. 2) *Our method can exploit invariant patterns to consistently alleviate the harmful effects of variant patterns under different distribution shift levels.* As the distribution shift level increases, almost all methods decline in performance since the relationship between variant patterns and labels goes stronger, so that the variant patterns are much easier to be exploited by the model, misleading the training process. However, the performance drop of **SILD** is significantly lower than baselines, which demonstrates that our method can alleviate the harmful effects of variant patterns under distribution shifts by exploiting invariant patterns in the spectral domain.

5.3 Ablation Studies

We conduct ablation studies to verify the effectiveness of the proposed disentangled spectrum mask and invariant spectral filtering in **SILD**. The ablated version ‘SILD w/o I’ removes invariant spectral filtering in **SILD** by setting $\lambda = 0$, and ‘SILD w/o M’ is trained without the disentangled spectrum masks. From Figure 3, we have the following observations. First, our proposed **SILD** outperforms all the variants as well as the best-performed baseline on all datasets, demonstrating the effectiveness of




Figure 3: Results of ablation studies, where ‘w/o I’ removes invariant spectral filtering in **SILD**, ‘w/o M’ removes disentangled spectrum masks, and ‘Best baseline’ denotes the best-performed baseline on each dataset. The error bars report the standard deviations. (Best viewed in color)

each component of our proposed method. Second, ‘SILD w/o I’ and ‘SILD w/o M’ drop drastically in performance on all datasets compared to the full version, which verifies that our proposed disentangled spectrum mask and spectral invariant learning can help the model to focus on invariant patterns to make predictions and significantly improve the performance under distribution shifts.

6 Related Works

Dynamic Graph Neural Networks Dynamic graphs ubiquitously exist in real-world applications [45, 46, 47, 48, 49, 50, 51, 52, 53] such as event forecasting, recommendation, etc. In comparison with static graphs [54, 55, 56, 57, 58, 59], dynamic graphs contain rich temporal information. Considerable research attention has been devoted to dynamic graph neural networks (DyGNNs) [1, 2, 60] to model the complex graph dynamics that include structures and features evolving through time. Some works adopt GNN to aggregate neighborhood information for each graph snapshot, and then utilize a sequence module to model the temporal information [61, 62, 63, 35, 24]. Some others utilize time-encoding techniques to encode the temporal links into time-aware embeddings and adopt a GNN or memory module [64, 65, 66, 67] to process structural information. Some other related works leverage spectral graph neural networks [68], global graph framelet convolution [69], and graph wavelets [70] to obtain better dynamic graph representations. However, distribution shifts remain largely unexplored in dynamic graph neural networks literature. The sole prior work DIDA [7] handles spatial-temporal distribution shifts on dynamic graphs in the temporal domain. To the best of our knowledge, this is the first study of handling distribution shifts on dynamic graphs in the spectral domain.

Out-of-Distribution Generalization A significant proportion of existing machine learning methodologies operate on the assumption that training and testing data are independent and identically distributed (i.i.d.). However, this assumption may not always hold true, especially in the context of complex real-world scenarios [71], and the uncontrollable distribution shifts between training and testing data distribution may lead to a significant decline in the model performance. Out-of-Distribution (OOD) generalization problem has recently drawn great attention in various areas [72, 71, 73]. Some works handle structural distribution shifts on static graphs [74, 15, 16, 75, 5, 76, 77, 78, 79, 80, 81] and temporal distribution shifts on time-series data [11, 12, 6, 13, 14, 82]. However, how to handle distribution shifts on dynamic graphs in the spectral domain remains unexplored.

Spectral Methods in Neural Networks The applications of spectral methods in neural networks have been broadly explored in many areas, including static graph data [83, 84, 85, 86, 87], time-series data [68, 88, 89, 90, 91], etc., for their advantages of modeling global patterns, powerful expressiveness and interpretability [92, 28]. Some work [93] proposes to reconstruct the image in the spectral domain to obtain robust image representations. Some work [29] proposes to augment the image data by perturbing the amplitude information in the spectral domain. Some work [94] proposes a multiwavelet-based method for compressing operator kernels. However, these methods are not applicable to dynamic graphs, not to mention the more complex scenarios under distribution shifts.

7 Conclusion

In this paper, we propose a novel model named Spectral Invariant Learning for Dynamic Graphs under Distribution Shifts (**SILD**), which can handle distribution shifts on dynamic graphs in the spectral domain. We design a DyGNN with Fourier transform to obtain the ego-graph trajectory spectrums. Then we propose a disentangled spectrum mask and invariant spectral filtering to discover the invariant and variant patterns in the spectral domain, and help the model rely on invariant spectral patterns to make predictions. Extensive experimental results on several synthetic and real-world datasets, including both node classification and link prediction tasks, demonstrate the superior performance of our method compared to state-of-the-art baselines under distribution shifts. One limitation is that in this paper we mainly focus on dynamic graphs in scenarios of discrete snapshots, and we leave extending our methods to continuous dynamic graphs for further explorations.

Acknowledgements

This work was supported by the National Key Research and Development Program of China No. 2020AAA0106300, National Natural Science Foundation of China (No. 62222209, 62250008, 62102222, 62206149), Beijing National Research Center for Information Science and Technology under Grant No. BNR2023RC01003, BNR2023TD03006, China National Postdoctoral Program for Innovative Talents No. BX20220185, China Postdoctoral Science Foundation No. 2022M711813, and Beijing Key Lab of Networked Multimedia. All opinions, findings, conclusions, and recommendations in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [1] Joakim Skarding, Bogdan Gabrys, and Katarzyna Musial. Foundations and modeling of dynamic networks using dynamic graph neural networks: A survey. *IEEE Access*, 9:79143–79168, 2021.
- [2] Yuecai Zhu, Fuyuan Lyu, Chengming Hu, Xi Chen, and Xue Liu. Learnable encoder-decoder architecture for dynamic graph: A survey. *arXiv preprint arXiv:2203.10480*, 2022.
- [3] Stephen J Brown, William Goetzmann, Roger G Ibbotson, and Stephen A Ross. Survivorship bias in performance studies. *The Review of Financial Studies*, 5(4):553–580, 1992.
- [4] Richard A Berk. An introduction to sample selection bias in sociological data. *American sociological review*, pages 386–398, 1983.
- [5] Qi Zhu, Natalia Ponomareva, Jiawei Han, and Bryan Perozzi. Shift-robust gnns: Overcoming the limitations of localized graph training data. *Advances in Neural Information Processing Systems*, 34, 2021.
- [6] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2021.
- [7] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Zhou Qin, and Wenwu Zhu. Dynamic graph neural networks under spatio-temporal distribution shift. In *Advances in Neural Information Processing Systems*, 2022.
- [8] Jayanta Mondal and Amol Deshpande. Managing large dynamic graphs efficiently. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 145–156, 2012.
- [9] Tuan Nhon Dang, Nick Pendar, and Angus Graeme Forbes. Timearcs: Visualizing fluctuations in dynamic networks. In *Computer Graphics Forum*, volume 35, pages 61–69. Wiley Online Library, 2016.
- [10] Jozef Barunik, Michael Ellington, et al. Dynamic networks in large financial and economic systems. *arXiv preprint arXiv:2007.07842*, 2020.
- [11] Jean-Christophe Gagnon-Audet, Kartik Ahuja, Mohammad-Javad Darvishi-Bayazi, Guillaume Dumas, and Irina Rish. Woods: Benchmarks for out-of-distribution generalization in time series tasks. *arXiv preprint arXiv:2203.09978*, 2022.
- [12] Yuntao Du, Jindong Wang, Wenjie Feng, Sinno Pan, Tao Qin, Renjun Xu, and Chongjun Wang. Adarnn: Adaptive learning and forecasting of time series. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 402–411, 2021.

- [13] Praveen Venkateswaran, Vinod Muthusamy, Vatche Isahagian, and Nalini Venkatasubramanian. Environment agnostic invariant risk minimization for classification of sequential datasets. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1615–1624, 2021.
- [14] Wang Lu, Jindong Wang, Yiqiang Chen, and Xinwei Sun. Diversify to generalize: Learning generalized representations for time series classification. *arXiv preprint arXiv:2209.07027*, 2021.
- [15] Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Discovering invariant rationales for graph neural networks. *arXiv preprint arXiv:2201.12872*, 2022.
- [16] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. Handling distribution shifts on graphs: An invariance perspective. *arXiv preprint arXiv:2202.02466*, 2022.
- [17] Mucong Ding, Kezhi Kong, Jiuhai Chen, John Kirchenbauer, Micah Goldblum, David Wipf, Furong Huang, and Tom Goldstein. A closer look at distribution shifts and out-of-distribution generalization on graphs. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.
- [18] Tian Jin, Qiong Wu, Xuan Ou, and Jianjun Yu. Community detection and co-author recommendation in co-author networks. *International Journal of Machine Learning and Cybernetics*, 12(2):597–609, 2021.
- [19] Wenjie Wang, Xinyu Lin, Fuli Feng, Xiangnan He, Min Lin, and Tat-Seng Chua. Causal representation learning for out-of-distribution recommendation. In *Proceedings of the ACM Web Conference 2022*, pages 3562–3571, 2022.
- [20] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [21] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. Invariant rationalization. In *International Conference on Machine Learning*, pages 1448–1458. PMLR, 2020.
- [22] Kartik Ahuja, Karthikeyan Shanmugam, Kush Varshney, and Amit Dhurandhar. Invariant risk minimization games. In *International Conference on Machine Learning*, pages 145–155. PMLR, 2020.
- [23] Jiaxuan You, Tianyu Du, and Jure Leskovec. Roland: graph learning framework for dynamic graphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2358–2366, 2022.
- [24] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. Dysat: Deep neural representation learning on dynamic graphs via self-attention networks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 519–527, 2020.
- [25] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. Evolvegnn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5363–5370, 2020.
- [26] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [27] Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
- [28] Kun Yi, Qi Zhang, Shoujin Wang, Hui He, Guodong Long, and Zhendong Niu. Neural time series analysis with fourier transform: A survey. *arXiv preprint arXiv:2302.02173*, 2023.
- [29] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14383–14392, 2021.
- [30] Bruce C Hansen and Robert F Hess. Structural sparseness and spatial phase alignment in natural scenes. *JOSA A*, 24(7):1873–1885, 2007.
- [31] Alan V Oppenheim and Jae S Lim. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981.
- [32] Leon N Piotrowski and Fergus W Campbell. A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, 11(3):337–346, 1982.

- [33] A Oppenheim, Jae Lim, Gary Kopec, and SC Pohlig. Phase in speech and pictures. In *ICASSP'79. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pages 632–637. IEEE, 1979.
- [34] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32, 2019.
- [35] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *International Conference on Neural Information Processing*, pages 362–373. Springer, 2018.
- [36] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- [37] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *EMNLP*, 2014.
- [38] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [39] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019.
- [40] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, pages 5815–5826. PMLR, 2021.
- [41] Jie Tang, Sen Wu, Jimeng Sun, and Hang Su. Cross-domain collaboration recommendation. In *KDD'2012*, 2012.
- [42] Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. Arnetminer: extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 990–998, 2008.
- [43] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-june Paul Hsu, and Kuansan Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, pages 243–246. ACM, 2015.
- [44] Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.
- [45] Lei Cai, Zhengzhang Chen, Chen Luo, Jiaping Gui, Jingchao Ni, Ding Li, and Haifeng Chen. Structural temporal graph neural networks for anomaly detection in dynamic graphs. In *Proceedings of the 30th ACM international conference on Information & Knowledge Management*, pages 3747–3756, 2021.
- [46] Songgaojun Deng, Huzefa Rangwala, and Yue Ning. Dynamic knowledge graph based multi-event forecasting. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1585–1595, 2020.
- [47] Jiaxuan You, Yichen Wang, Aditya Pal, Pong Eksombatchai, Chuck Rosenburg, and Jure Leskovec. Hierarchical temporal convolutional networks for dynamic recommender systems. In *The world wide web conference*, pages 2236–2246, 2019.
- [48] Yanbang Wang, Pan Li, Chongyang Bai, and Jure Leskovec. Tedic: Neural modeling of behavioral patterns in dynamic social interaction networks. In *Proceedings of the Web Conference 2021*, pages 693–705, 2021.
- [49] Haoyang Li, Peng Cui, Chengxi Zang, Tianyang Zhang, Wenwu Zhu, and Yishi Lin. Fates of microscopic social ecosystems: Keep alive or dead? In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 668–676, 2019.
- [50] Jiapeng Wu, Meng Cao, Jackie Chi Kit Cheung, and William L Hamilton. Temp: Temporal message passing for temporal knowledge graph completion. *arXiv preprint arXiv:2010.03526*, 2020.
- [51] Zeyang Zhang, Ziwei Zhang, Xin Wang, Yijian Qin, Zhou Qin, and Wenwu Zhu. Dynamic heterogeneous graph attention neural architecture search. In *Thirty-Seventh AAAI Conference on Artificial Intelligence*, 2023.

- [52] Zeyang Zhang, Xingwang Li, Fei Teng, Ning Lin, Xueling Zhu, Xin Wang, and Wenwu Zhu. Out-of-distribution generalized dynamic graph neural network for human albumin prediction. In *IEEE International Conference on Medical Artificial Intelligence*, 2023.
- [53] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Yijian Qin, Simin Wu, and Wenwu Zhu. Llm4dyg: Can large language models solve problems on dynamic graphs? *arXiv preprint*, 2023.
- [54] Ziwei Zhang, Peng Cui, and Wenwu Zhu. Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 34(1):249–270, 2020.
- [55] Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Guangyao Shen, Shiqi Shen, and Wenwu Zhu. Unsupervised graph neural architecture search with disentangled self-supervision. In *Advances in Neural Information Processing Systems*, 2023.
- [56] Chaoyu Guan, Ziwei Zhang, Haoyang Li, Heng Chang, Zeyang Zhang, Yijian Qin, Jiyan Jiang, Xin Wang, and Wenwu Zhu. Autogl: A library for automated graph learning. In *ICLR 2021 Workshop GTRL*, 2021.
- [57] Yijian Qin, Ziwei Zhang, Xin Wang, Zeyang Zhang, and Wenwu Zhu. Nas-bench-graph: Benchmarking graph neural architecture search. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022.
- [58] Yijian Qin, Xin Wang, Zeyang Zhang, and Wenwu Zhu. Graph differentiable architecture search with structure learning. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [59] Ziwei Zhang, Haoyang Li, Zeyang Zhang, Yijian Qin, Xin Wang, and Wenwu Zhu. Large graph models: A perspective. *arXiv preprint arXiv:2308.14522*, 2023.
- [60] Chao Chen, Haoyu Geng, Nianzu Yang, Xiaokang Yang, and Junchi Yan. Easydgl: Encode, train and interpret for continuous-time dynamic graph learning. *arXiv preprint arXiv:2303.12341*, 2023.
- [61] Menglin Yang, Min Zhou, Marcus Kalander, Zengfeng Huang, and Irwin King. Discrete-time temporal network embedding via implicit hierarchical learning in hyperbolic space. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1975–1985, 2021.
- [62] Li Sun, Zhongbao Zhang, Jiawei Zhang, Feiyang Wang, Hao Peng, Sen Su, and Philip S Yu. Hyperbolic variational graph neural network for modeling dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4375–4383, 2021.
- [63] Ehsan Hajiramezanali, Arman Hasanzadeh, Krishna Narayanan, Nick Duffield, Mingyuan Zhou, and Xiaoning Qian. Variational graph recurrent neural networks. *Advances in neural information processing systems*, 32, 2019.
- [64] Yanbang Wang, Yen-Yu Chang, Yunyu Liu, Jure Leskovec, and Pan Li. Inductive representation learning in temporal networks via causal anonymous walks. *arXiv preprint arXiv:2101.05974*, 2021.
- [65] Weilin Cong, Yanhong Wu, Yuandong Tian, Mengting Gu, Yinglong Xia, Mehrdad Mahdavi, and Chuncheng Jason Chen. Dynamic graph representation learning via graph transformer networks. *arXiv preprint arXiv:2111.10447*, 2021.
- [66] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. *arXiv preprint arXiv:2002.07962*, 2020.
- [67] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637*, 2020.
- [68] Defu Cao, Yujing Wang, Juanyong Duan, Ce Zhang, Xia Zhu, Congrui Huang, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, et al. Spectral temporal graph neural network for multivariate time-series forecasting. *Advances in neural information processing systems*, 33:17766–17778, 2020.
- [69] Bingxin Zhou, Xinliang Liu, Yuehua Liu, Yunying Huang, Pietro Lio, and Yu Guang Wang. Well-conditioned spectral transforms for dynamic graph representation. In *Learning on Graphs Conference*, pages 12–1. PMLR, 2022.
- [70] Anson Bastos, Abhishek Nadgeri, Kuldeep Singh, Toyotaro Suzumura, and Manish Singh. Learnable spectral wavelets on dynamic graphs to capture global interactions. *arXiv preprint arXiv:2211.11979*, 2022.

- [71] Zheyang Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*, 2021.
- [72] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *arXiv e-prints*, pages arXiv–2103, 2021.
- [73] Huaxiu Yao, Yu Wang, Sai Li, Linjun Zhang, Weixin Liang, James Zou, and Chelsea Finn. Improving out-of-distribution robustness via selective augmentation. In *Proceeding of the Thirty-ninth International Conference on Machine Learning*, 2022.
- [74] Haoyang Li, Xin Wang, Ziwei Zhang, and Wenwu Zhu. Out-of-distribution generalization on graphs: A survey. *arXiv preprint arXiv:2202.07987*, 2022.
- [75] Yongqiang Chen, Yonggang Zhang, Han Yang, Kaili Ma, Binghui Xie, Tongliang Liu, Bo Han, and James Cheng. Invariance principle meets out-of-distribution generalization on graphs. *arXiv preprint arXiv:2202.05441*, 2022.
- [76] Yijian Qin, Xin Wang, Ziwei Zhang, Pengtao Xie, and Wenwu Zhu. Graph neural architecture search under distribution shifts. In *International Conference on Machine Learning*, pages 18083–18095. PMLR, 2022.
- [77] Haoyang Li, Xin Wang, Ziwei Zhang, and Wenwu Zhu. Ood-gnn: Out-of-distribution generalized graph neural network. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [78] Zeyang Zhang, Ziwei Zhang, Xin Wang, and Wenwu Zhu. Learning to solve travelling salesman problem with hardness-adaptive curriculum. *arXiv preprint arXiv:2204.03236*, 2022.
- [79] Ziwei Zhang, Xin Wang, Zeyang Zhang, Peng Cui, and Wenwu Zhu. Revisiting transformation invariant geometric deep learning: Are initial representations all you need? *arXiv preprint arXiv:2112.12345*, 2021.
- [80] Shaohua Fan, Xiao Wang, Chuan Shi, Peng Cui, and Bai Wang. Generalizing graph neural networks on out-of-distribution graphs. *arXiv preprint arXiv:2111.10657*, 2021.
- [81] Haoyang Li, Ziwei Zhang, Xin Wang, and Wenwu Zhu. Learning invariant graph representations for out-of-distribution generalization. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022.
- [82] Huaxiu Yao, Caroline Choi, Yoonho Lee, Pang Wei Koh, and Chelsea Finn. Wild-time: A benchmark of in-the-wild distribution shift over time. In *Proceedings of the Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [83] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.
- [84] Bingbing Xu, Huawei Shen, Qi Cao, Yunqi Qiu, and Xueqi Cheng. Graph wavelet neural network. *arXiv preprint arXiv:1904.07785*, 2019.
- [85] Henry Kenlay, Dorina Thanou, and Xiaowen Dong. Interpretable stability bounds for spectral graph filters. In *International conference on machine learning*, pages 5388–5397. PMLR, 2021.
- [86] Xiyuan Wang and Muhan Zhang. How powerful are spectral graph neural networks. In *International Conference on Machine Learning*, pages 23341–23362. PMLR, 2022.
- [87] Deyu Bo, Chuan Shi, Lele Wang, and Renjie Liao. Specformer: Spectral graph neural networks meet transformers. *arXiv preprint arXiv:2303.01028*, 2023.
- [88] Robert Salomone, Matias Quiroz, Robert Kohn, Mattias Villani, and Minh-Ngoc Tran. Spectral subsampling mcmc for stationary time series. In *International Conference on Machine Learning*, pages 8449–8458. PMLR, 2020.
- [89] Henning Lange, Steven L Brunton, and J Nathan Kutz. From fourier to koopman: Spectral methods for long-term time series prediction. *The Journal of Machine Learning Research*, 22(1):1881–1918, 2021.
- [90] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning*, pages 27268–27286. PMLR, 2022.
- [91] Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. Self-supervised contrastive pre-training for time series via time-frequency consistency. *Advances in Neural Information Processing Systems*, 35:3988–4003, 2022.

- [92] Deyu Bo, Xiao Wang, Yang Liu, Yuan Fang, Yawen Li, and Chuan Shi. A survey on spectral graph neural networks. *arXiv preprint arXiv:2302.05631*, 2023.
- [93] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13919–13929, 2021.
- [94] Gaurav Gupta, Xiongye Xiao, and Paul Bogdan. Multiwavelet-based operator learning for differential equations. *Advances in neural information processing systems*, 34:24048–24062, 2021.