

Handling Missing Data via Max-Entropy Regularized Graph Autoencoder

Ziqi Gao², Yifan Niu¹, Jiashun Cheng², Jianheng Tang², Lanqing Li^{3*},
Tingyang Xu³, Peilin Zhao³, Fugee Tsung^{1,2}, Jia Li^{1,2*}

¹The Hong Kong University of Science and Technology (Guangzhou)

²The Hong Kong University of Science and Technology

³AI Lab, Tencent

zgaoat@connect.ust.hk, yniu669@connect.hkust-gz.edu.cn,
{jchengak, jtangbf}@connect.ust.hk, {tingyangxu, masonzhao, lanqingli}@tencent.com,
{season, jialeel}@ust.hk

Abstract

Graph neural networks (GNNs) are popular weapons for modeling relational data. Existing GNNs are not specified for attribute-incomplete graphs, making missing attribute imputation a burning issue. Until recently, many works notice that GNNs are coupled with spectral concentration, which means the spectrum obtained by GNNs concentrates on a local part in spectral domain, e.g., low-frequency due to oversmoothing issue. As a consequence, GNNs may be seriously flawed for reconstructing graph attributes as graph spectral concentration tends to cause a low imputation precision. In this work, we present a regularized graph autoencoder for graph attribute imputation, named MEGAE, which aims at mitigating spectral concentration problem by maximizing the graph spectral entropy. Notably, we first present the method for estimating graph spectral entropy without the eigen-decomposition of Laplacian matrix and provide the theoretical upper error bound. A maximum entropy regularization then acts in the latent space, which directly increases the graph spectral entropy. Extensive experiments show that MEGAE outperforms all the other state-of-the-art imputation methods on a variety of benchmark datasets.

1 Introduction

Graph attribute missing is ubiquitous due to messy collection and recording process (Yoon, Davtyan, and van der Schaar 2016; Roth 1994). In a biochemical scenario, for instance, data is missed due to the difficulties in measuring or calculating quantitative molecular properties (e.g., HOMO and LUMO orbital energy) at the atomic level (Yomogida et al. 2012; Salomon et al. 2003), attenuating the graph or node representations when we introduce a graph model for molecules. For graph learning tasks, even a small fraction of missing attributes will potentially interfere with the performance, leading to biased inferences (Sterne et al. 2009; Wooldridge 2007).

Graph Neural Networks (GNNs) have demonstrated powerful ability on various graph-related tasks. Most GNNs consist of two parts: feature propagation and transformation. The former is the key to the success of GNNs as it can yield suitable spectral responses for graph desires (Li et al. 2022). For example, graph convolutional networks (GCN) (Kipf and Welling

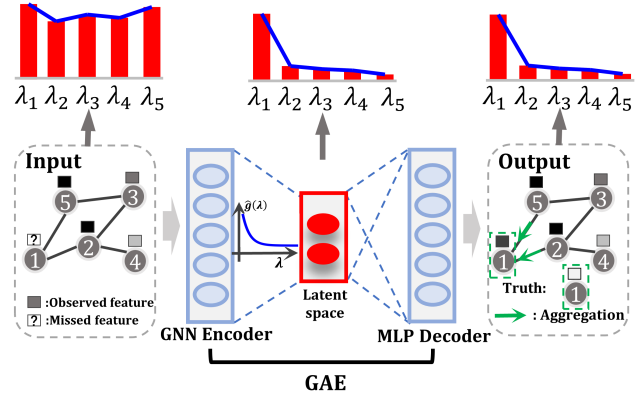


Figure 1: Spectral visualization of GAE. Red histograms show the graph spectra. After GAE encoder (GCN), spectral concentration appears in the latent space. Upon MLP decoder, the spectral concentration is not alleviated.

2017) applies a recursive aggregation mechanism for updating node representations, leading to the low-frequency concentration from a spectral view (Wu et al. 2019; Donnat et al. 2018; Tang et al. 2022). The mechanism makes GCN effective for numerous real-world applications as low-frequency concentration promotes the convergence of representations of adjacent nodes (Nt and Maehara 2019; Xu et al. 2020). However, this spectral concentration diminishes partial eigenvalues, i.e., spectral values. As a consequence, it causes serious performance degradation on the standard graph attribute imputation models such as graph autoencoders (GAEs) (Li et al. 2020; Park et al. 2019; Li et al. 2021). In Figure 1, we visualize the graph spectra of the input, latent and output based on GAE (Kipf and Welling 2016), which is the backbone of most GAEs. As the figure makes clear, existing GAEs leave sufficient room for improvement to combat the loss of spectral components. Delving into graph attribute reconstruction, a natural question arises: *how can GAEs impute with alleviated spectral concentration?*

In this work, inspired by the maximum entropy principle (Narayan and Nityananda 1986; Burch, Gull, and Skilling 1983; Frieden 1972) that is capable to alleviate concentrated data distributions, we propose to *maximize the graph spectral entropy* for mitigating spectral concentration while imputa-

*Corresponding author.

tion. We present a novel method for graph attribute imputation called Max-Entropy Graph AutoEncoder (MEGAE). Unlike existing imputation methods, our model trains a GAE with the regularization for maximizing the graph spectral entropy. However, when it comes to the computation of graph spectral entropy, it becomes challenging due to large time complexity of the eigen-decomposition for Laplacian matrix (Cai, He, and Han 2007a,b; Jia et al. 2016). In this regard, we design the tight wavelet kernels (Shuman et al. 2015; Leonardi and Van De Ville 2013; Ahmad and Sheikh 2018) in MEGAE to encode features into the spectral domain involving several spectral bands. We then theoretically prove that the proposed wavelet paradigm can well approximate the actual graph spectral entropy by a predictable upper error bound without eigen-decomposition. Upon entropy regularized latent representations, we use symmetric wavelet decoder to perform data deconvolution for reconstructing graph attributes.

In the experiments, we proceed with an empirical evaluation of MEGAE on single- and multi-graph datasets. Firstly, we show that MEGAE outperforms 16 state-of-the-art methods on attribute imputation tasks, including commonly used classical methods and graph learning based imputation methods. Additional downstream experiments of graph and node classifications demonstrate that imputed graphs by MEGAE can obtain the best accuracy performance when compared to state-of-the-art imputation models.

Contributions. We summarize our contributions below:

- We propose to maximize the graph spectral entropy to overcome the spectral concentration issue in GAEs for graph attribute reconstruction.
- We present Max-Entropy Graph AutoEncoder (MEGAE) for encoding graphs into spectral domain and maximizing the spectral entropy in the latent space.
- We develop an efficient method for approximating the actual graph spectral entropy without eigen-decomposition, and importantly, provide the theoretical upper error bound.

2 Related Work

Matrix Completion. Missing data is a widely researched topic. Matrix completion methods can be applied to impute graph attributes without using the graph structures. Most proposed methods impute with a joint distribution on the incomplete data. For instance, the joint modeling methods impute by drawing from the predictive distribution, including Bayesian strategies (Murray and Reiter 2016), matrix completion methods (Candès and Recht 2009; Hastie et al. 2015), and Generative Adversarial Networks (Yoon, Jordon, and Schaar 2018; Yoon and Sull 2020). Another way of joint modeling involves iteratively imputing values of each variable using chained equations (Van Buuren and Groothuis-Oudshoorn 2011) formulated with other variables (White, Royston, and Wood 2011; Van Buuren 2018; Muzellec et al. 2020). Discriminative models such as random forests (Xia et al. 2017), distribution constraints using optimal transport (Muzellec et al. 2020) and causally-aware

imputation (Kyono et al. 2021) tend to depend on strong assumptions, which may result in a lack of flexibility to handle mixed-mode data. Most importantly, as for graph scenarios, matrix completion methods are generally limited without awareness of the underlying graph structures.

Attribute Imputation with Graph Learning. Recently, graph learning models have been used to tackle the imputation task. GC-MC (Berg, Kipf, and Welling 2017) and IGM (Zhang and Chen 2019) construct helpful bipartite interaction graphs to impute with a given adjacency matrix as side information. Then a GAE is applied to predict the absent features. Gaussian mixture model is utilized for imputation stability under a high missing rate (Monti et al. 2017). GRAPE (You et al. 2020) combines imputation and representation learning, which can well impute features of continuous variables but tend not to perform well on datasets containing nodes with all features missed (Tu et al. 2021). GDN (Li et al. 2021) can impute from over-smoothed representations with a given graph structure. In this case, relationships between nodes are explicitly encoded for further imputation. Above methods show that graph learning is competent for feature reconstruction with structure dependencies. Notably, only GRAPE, GDN and our method are applicable to both discrete and continuous features. Despite both GDN and our method focus on mechanistic improvements for existing GAEs, we explore GAEs with a more general phenomenon of concentration in a spectral view and not just for recovering high-frequency details.

3 Preliminary

In this section, we provide the formulation of graph attribute imputation problem in Section 3.1. Then we introduce the regularized object, named graph spectral entropy in Section 3.2. In Section 3.3, we describe wavelet entropy, a core concept for efficiently approximating the graph spectral entropy.

3.1 Problem Formulation

For an undirected graph $G = (\mathbf{A}, \mathbf{X})$, $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix and $\mathbf{X} \in \mathbb{R}^{N \times D}$ is a complete feature matrix where \mathbf{X}_{ij} denotes the graph attribute of i -th node in the j -th feature dimension. In the problem of graph attribute imputation, $\mathbf{R} \in \{0, 1\}^{N \times D}$ is defined as the mask matrix whose element $\mathbf{R}_{ij} = 1$ if \mathbf{X}_{ij} is observed and $\mathbf{R}_{ij} = 0$ otherwise.

The objective of this work is to predict the missing graph attributes \mathbf{X}_{ij} at $\mathbf{R}_{ij} = 0$. Formally, we aim to develop a mapping $f(\cdot)$ to generate the imputed data matrix $\tilde{\mathbf{X}} \in \mathbb{R}^{N \times D}$ defined as

$$\tilde{\mathbf{X}} = f(\mathbf{X}, \mathbf{R}, \mathbf{A}). \quad (1)$$

3.2 Graph Spectral Entropy

The *graph Laplacian* is defined as $\mathbf{L} = \mathbf{I} - \tilde{\mathbf{A}}$, where $\tilde{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ denotes the normalized adjacency matrix and $\mathbf{D} = \text{diag}(\sum_k \mathbf{A}_{1k}, \dots, \mathbf{A}_{Nk})$ denotes the diagonal degree matrix. With eigen-decomposition, \mathbf{L} can be decomposed into $\mathbf{L} = \mathbf{U}^T \mathbf{\Lambda} \mathbf{U}$, where \mathbf{U} consists of the eigenvectors of \mathbf{L} . $\mathbf{\Lambda} = \text{diag}(\lambda)$ is the diagonal matrix whose diagonal

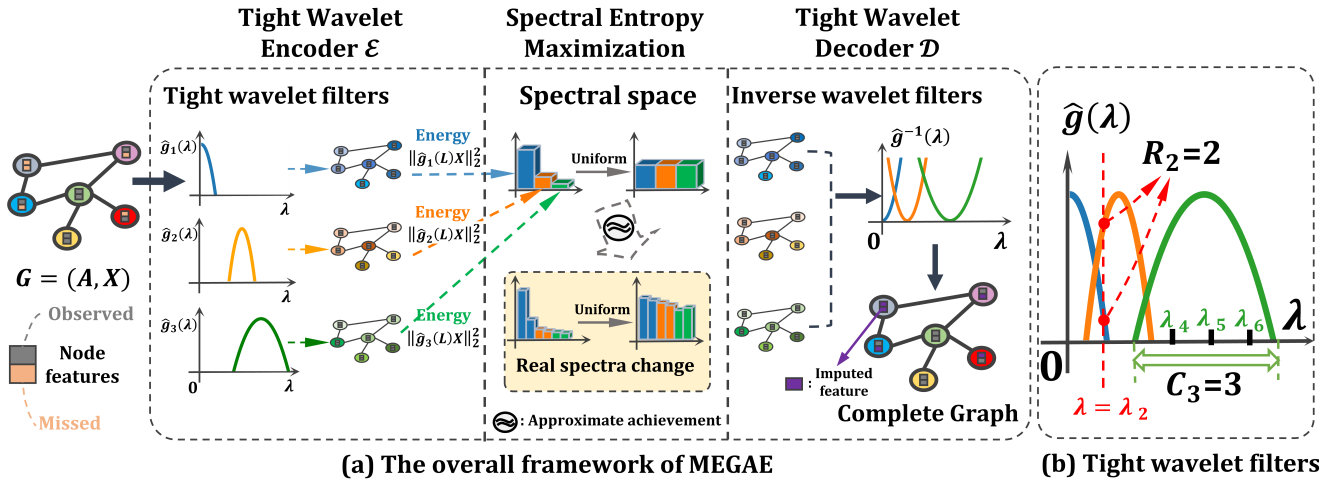


Figure 2: (a) The proposed MEGAE for graph attribute imputation. With tight wavelet filters, the encoder filters the input graph into $\hat{g}_1(L)X, \dots, \hat{g}_M(L)X$ ($M = 3$ here), of which the energies (i.e., blue, orange and green bars) are obtained with the L-2 norm function. We maximize wavelet entropy of M energies, which are encouraged to be more uniform in distributions. This wavelet entropy maximization induces the same tendency to maximization for the real graph spectral entropy. We apply inverse wavelet filters to decode the data from the latent (spectral) space for final reconstruction. (b) A diagram to explain the Coverage and Crossness of given tight wavelet filters. We provide related descriptions in Section 4.1.

elements $\lambda = \{\lambda_1, \dots, \lambda_N\}$ are eigenvalues of L . Consider a feature vector $\mathbf{x} = \{x_1, \dots, x_N\}^\top \in \mathbb{R}^N$ as a single feature dimension of the whole matrix \mathbf{X} , the Fourier transform $\hat{\mathbf{x}} = \mathbf{U}^\top \mathbf{x} = \{\hat{x}_1, \dots, \hat{x}_N\}^\top$ would be the most proven method for representing spectrum, of which the entropy is defined as *graph spectral entropy* in Definition 3.1.

Definition 3.1 (Graph Spectral Entropy). *Given the feature vector \mathbf{x} , its spectral energy at λ_i is denoted as \hat{x}_i^2 and total energy is $E_s = \sum_{i=1}^N \hat{x}_i^2 = \|\mathbf{U}^\top \mathbf{x}\|_2^2$. Then, the graph spectral entropy of \mathbf{x} is defined as*

$$\xi_s(\mathbf{x}, L) = - \sum_{i=1}^N \frac{\hat{x}_i^2}{E_s} \log \frac{\hat{x}_i^2}{E_s}.$$

Intuitively, maximizing graph spectral entropy would ensure the spectral distribution to be relatively more uniform and keep information in any spectral pass-band, as opposed to spectral concentration in existing GAEs (Park et al. 2019; Wu et al. 2019). However, one notable issue with the computation of graph spectral entropy is that eigen-decomposition is indispensable, which largely limits the scalability for computing on large-scale data matrix.

3.3 Wavelet Entropy

We introduce wavelet theory for approximating the graph spectral entropy with wavelet transform. Unlike Fourier transform that we formulate in section 3.2, wavelet transform (Xu et al. 2019; Hammond, Vandergheynst, and Gribonval 2011) is able to *cover ranges of spectra via different pass-band kernels* $\hat{g} = \{\hat{g}_1, \dots, \hat{g}_M\}$. It shows the possibility of representing graph spectral entropy by wavelet entropy with well-calibrated kernels.

The wavelet transform employs a set of M wavelets as bases, defined as $\mathcal{W} = \{\mathcal{W}_{\hat{g}_1}, \dots, \mathcal{W}_{\hat{g}_M}\}$. Mathematically, the

wavelet transform of \mathbf{x} is defined as

$$\mathcal{W}_{\hat{g}_m}(\mathbf{x}) = \mathbf{U} \hat{g}_m(\Lambda) \mathbf{U}^\top \mathbf{x}, \quad (2)$$

where $\hat{g}_m(\cdot)$ is a wavelet kernel function defined in spectral domain. For computational efficiency, in many cases, kernel function \hat{g}_m should be smooth and easily approximated by polynomials such that $\mathbf{U} \hat{g}_m(\Lambda) \mathbf{U}^\top = \hat{g}_m(L)$.

Definition 3.2 (Wavelet Entropy). *Given the feature vector \mathbf{x} and a set of wavelet kernel functions $\hat{g} = \{\hat{g}_1, \dots, \hat{g}_M\}$, the total wavelet energy of \mathbf{x} is $E_w = \sum_{m=1}^M \|\mathcal{W}_{\hat{g}_m}(\mathbf{x})\|_2^2$. Then, the wavelet entropy of \mathbf{x} is defined as*

$$\xi_w(\mathbf{x}, L, \hat{g}) = - \sum_{m=1}^M \frac{\|\mathcal{W}_{\hat{g}_m}(\mathbf{x})\|_2^2}{E_w} \log \frac{\|\mathcal{W}_{\hat{g}_m}(\mathbf{x})\|_2^2}{E_w}.$$

4 Maximizing Graph Spectral Entropy

This section describes our main contributions. Following our motivation to efficiently maximize the graph spectral entropy while imputing missing attributes, we propose an alternative method for approximating the graph spectral entropy *without eigen-decomposition* and provide the upper bound for approximation error. On this basis, we propose our model called Max-Entropy Graph AutoEncoder (MEGAE).

4.1 Graph Wavelet for Approximation

In this section, we present our main idea that graph spectral entropy can be well approximated by wavelet entropy. To achieve the goal, the crucial issue is how to construct a rough spectral space for approximation with the efficient graph convolutions.

Definition 4.1 (Tight Wavelet Frame on Graph). *A set of wavelet kernel functions $\hat{g} = \{\hat{g}_1, \dots, \hat{g}_M\}$ forms a tight*

frame on graph if

$$G(\lambda_i) := \sum_{m=1}^M [\hat{g}_m(\lambda_i)]^2 = 1, \forall \lambda_i \in \lambda.$$

Inspired by the fact that tight wavelet frame can ensure energy preservation (Definition 4.1) in graph signal processing (Leonardi and Van De Ville 2013), we can smoothly deduce a key connection between wavelet energy (formulated in Section 3.3) and graph spectral energy (formulated in Section 3.2), which are both key components in the respective wavelet and spectral entropy. Proposition 4.2 describes this connection, indicating that the total wavelet energy is strictly equivalent to the total graph spectral energy. The proof is trivial and illustrated in Appendix A.

Proposition 4.2 (Energy Parseval’s Identity). *If \hat{g} is a tight wavelet frame, the total wavelet energy E_w is equivalent to total graph spectral energy E_s . Thus, the Parseval’s identity holds for the spectral domain, that is*

$$E_w = E_s.$$

The wavelet energy consists of M energy components and each of them is derived from the filtered graph features by a wavelet kernel. Therefore, intuitively, a more distinguishable kernel function tends to provide accurate approximation of graph spectral entropy. To quantify the kernel distinguishability of tight wavelet, we define Coverage C_m and Crossness R_i . Formally, C_m is the number of spectrum covered by kernel function \hat{g}_m denoted as

$$C_m = |\{i | \hat{g}_m(\lambda_i) \neq 0, i = 1, \dots, N\}|, \quad (3)$$

and R_i is the number of activated kernel functions on λ_i represented by

$$R_i = |\{m | \hat{g}_m(\lambda_i) \neq 0, m = 1, \dots, M\}|. \quad (4)$$

A toy example is provided in Figure 2(b) for easier understanding. Coverage C_3 represents the number of eigenvalues (i.e., $\lambda_4, \lambda_5, \lambda_6$) lying within the domain of 3rd wavelet kernel (i.e., the green curve). The Crossness R_2 represents number of intersection (red) points of all three wavelet kernels and line $\lambda = \lambda_2$.

Proposition 4.3 (Approximation of graph spectral entropy). *If \hat{g} is a tight frame, the approximation error between graph spectral entropy $\xi_s(\mathbf{x}, \mathbf{L})$ and wavelet entropy $\xi_w(\mathbf{x}, \mathbf{L}, \hat{g})$ is bounded by*

$$|\xi_s(\mathbf{x}, \mathbf{L}) - \xi_w(\mathbf{x}, \mathbf{L}, \hat{g})| \leq e(\hat{g}, \lambda)$$

where $e(\hat{g}, \lambda) = \max(\{\log C_m\}_{m=1}^M \cup \{\log R_i\}_{i=1}^N)$.

Please see Appendix B for the proof. To better understand Proposition 4.3, let $N = M$ and all kernels be disjoint, which means the number of wavelet kernels equals to the number of nodes in the graph. In this perfect situation, one kernel function only covers one eigenvalue (i.e., Coverage $C_m = 1$) and one eigenvalue is only covered by one kernel function (i.e., Crossness $R_m = 1$). Then, the approximation error is zero.

For applying a tight wavelet frame with high distinguishability, we follow the mechanisms in (Shuman et al. 2015) to design a set of wavelet kernel function $g = \{\hat{g}_1, \dots, \hat{g}_M\}$. The design details are in Appendix C.

4.2 The Proposed MEGAE

Based on the designed tight wavelet kernels, we propose MEGAE, a deterministic graph autoencoder framework for graph missing attribute imputation. As demonstrated in Figure 2, MEGAE consists of two key modules: (1) wavelet based graph autoencoder and (2) maximum entropy regularization.

Encoder. With the definitions of an undirected graph provided in Section 3, we employ M tight wavelet kernels described in Section 4.1 in parallel to perform graph wavelet convolution. Specifically, the convolutional layer of m -th channel is expressed as

$$\mathcal{E}_1(\mathbf{Z}_m^{(0)}, \mathbf{L}, \hat{g}_m) = \phi(\mathcal{W}_{\hat{g}_m}(\mathbf{Z}_m^{(0)})\mathbf{W}_m^{(0)}), \quad (5)$$

where $\mathbf{Z}_m^{(0)} = \mathbf{X} \odot \mathbf{R}$ and $\phi(\cdot)$ is the activation function such as leaky ReLU. To admit a fast algorithm for wavelet transform, Maclaurin series approximation (Li et al. 2021) of order K is applied such that $\mathcal{W}_{\hat{g}_m}(\mathbf{Z}_m^{(0)}) = \sum_{k=0}^K \alpha_{m,k} \mathbf{L}^k \mathbf{Z}_m^{(0)}$. We introduce one more layer of perceptron to enrich the representation in latent space, denoted as

$$\mathcal{E}_2(\mathbf{Z}_m^{(1)}) = \phi(\mathbf{Z}_m^{(1)}\mathbf{W}_m^{(1)}). \quad (6)$$

The *Encoder* is thus formulated as

$$\begin{aligned} \mathbf{Z}_m^{(1)} &= \mathcal{E}_1(\mathbf{Z}_m^{(0)}, \mathbf{L}, \hat{g}_m), \\ \mathbf{Z}_m^{(2)} &= \mathcal{E}_2(\mathbf{Z}_m^{(1)}). \end{aligned} \quad (7)$$

Decoder. Given the observed information, our decoder aims to impute missing values from latent space. As the inverse of graph wavelet convolution, the graph wavelet deconvolutional layer of m -th channel is expressed as

$$\mathcal{D}_1(\mathbf{Z}_m^{(2)}, \mathbf{L}, \hat{g}_m^{-1}) = \phi(\mathcal{W}_{\hat{g}_m^{-1}}(\mathbf{Z}_m^{(2)})\mathbf{W}_m^{(2)}), \quad (8)$$

where $\mathcal{W}_{\hat{g}_m^{-1}}(\mathbf{Z}_m^{(2)}) = \sum_{k=1}^K \beta_{m,k} \mathbf{L}^k \mathbf{Z}_m^{(2)}$ is provided by Maclaurin series. We then aggregate results generated by M channels to construct the imputation layer as

$$\mathcal{D}_2(\mathbf{Z}_1^{(3)}, \dots, \mathbf{Z}_M^{(3)}) = \phi(\mathbf{Z}_{AGG}^{(3)}\mathbf{W}^{(3)}), \quad (9)$$

where $\mathbf{Z}_{AGG}^{(3)} = \text{AGG}([\mathbf{Z}_1^{(3)}, \dots, \mathbf{Z}_M^{(3)}])$ and $\text{AGG}(\cdot)$ denotes aggregation function such as concatenation.

The *Decoder* is thus formulated as

$$\begin{aligned} \mathbf{Z}_m^{(3)} &= \mathcal{D}_1(\mathbf{Z}_m^{(2)}, \mathbf{L}, \hat{g}_m^{-1}), \\ \tilde{\mathbf{X}} &= \mathcal{D}_2(\mathbf{Z}_1^{(3)}, \dots, \mathbf{Z}_M^{(3)}). \end{aligned} \quad (10)$$

Optimization. Our network is jointly optimized with respect to two objectives. We start from the imputation precision, which is the main focus in missing data problem. Intuitively, the reconstruction loss is defined as

$$\mathcal{L}_R = \left\| (\tilde{\mathbf{X}} - \mathbf{X}) \odot (\mathbf{1}_{N \times D} - \mathbf{R}) \right\|_2. \quad (11)$$

Given our imputation motivation, MEGAE introduces spectral entropy regularization to aid data imputation. Based on the discussion in Section 4.1, graph spectral entropy can be well substituted by wavelet entropy. Let $\mathbf{Z}_m^{(2)}(:, d)$ be the d -th

Methods	ENZYMES	PRO_full	QM9	Synthie	FRANKE	FIRST_DB	ENZYMES	PRO_full
	RMSE with 0.1 missing features						ACC.	ACC.
MEAN	0.0602	0.0653	0.2983	0.2063	0.3891	0.1500	62.06%	69.96%
KNN	0.0350	0.0329	0.3058	0.1718	0.2010	0.1296	63.53%	72.23%
SVD	0.0783	0.0465	0.2524	0.1697	0.2766	0.1685	61.60%	71.16%
MICE	0.0292	0.0210	0.1986	0.1899	0.1359	0.1036	64.46%	73.50%
GAIN	0.0300	0.0245	0.1973	0.1649	0.1103	0.0909	64.42%	73.21%
OT	0.0323	0.0206	0.2003	0.1865	0.1116	0.0892	64.13%	73.62%
MIRACLE	0.0288	0.0188	0.1846	0.1632	0.1196	0.0889	65.03%	74.07%
GraphVAE	0.0357	0.0199	0.1579	0.1898	0.1099	0.1202	63.46%	73.86%
MolGAN	0.0326	0.0160	0.1478	0.1864	0.1078	0.1379	64.16%	74.20%
GRAPE	0.0302	0.0147	0.1869	0.1798	0.1069	0.0986	64.48%	74.53%
GDN	0.0267	0.0155	0.1598	0.1764	0.1066	0.0869	65.57%	74.50%
Inv_GCN	0.0256	0.0135	0.1773	0.1446	0.1089	0.1032	66.00%	74.98%
Inv_MLP	0.0254	0.0129	0.1499	0.1520	0.1064	0.0872	66.06%	75.07%
MEGAE	0.0223	0.0099	0.1396	0.1203	0.0936	0.0789	66.27%	75.97%
Performance	0.0044	0.0048	0.0082	0.0429	0.0128	0.0080	0.70%	1.44%
gain	0.0560	0.0554	0.1662	0.0860	0.2955	0.0896	4.67%	6.01%

Table 1: RMSE results on 6 multi-graph datasets and graph classification accuracy on two datasets. After running 5 trials, we report the mean results of which the best method is bolded and the second best is underlined. Performance gains indicate the maximum (lower) and minimum (upper) gains of state-of-the-art (MEGAE) compared to other baselines. Note that we abbreviate 'PROTEIN_full', 'FRANKENSTEIN' and 'FIRSTMM_DB' to 'PRO_full', 'FRANKE' and 'FIRST_DB', respectively.

dimension of $\mathbf{Z}_m^{(2)}$, the graph spectral entropy loss is thus defined as

$$\mathcal{L}_S = -\frac{1}{D} \sum_{d=1}^D \sum_{m=1}^M P_m^d \log P_m^d, \quad (12)$$

where

$$P_m^d = \frac{\left\| \mathbf{Z}_m^{(2)}(:, d) \right\|_2^2}{\sum_{m=1}^M \left\| \mathbf{Z}_m^{(2)}(:, d) \right\|_2^2}. \quad (13)$$

Finally, by combining the aforementioned losses, the overall objective function of MEGAE is formulated as

$$\mathcal{L} = \mathcal{L}_R - \gamma \mathcal{L}_S, \quad (14)$$

where γ is a hyperparameter to balance the regularization of graph spectral entropy and reconstruction loss.

5 Experiments

In this section, we validate the performance of MEGAE using a variety of datasets. We evaluate the effectiveness of MEGAE on two categories of graph datasets:

- **Type 1: Imputation on multi-graph datasets.** We impute the missing graph attributes on multi-graph datasets, e.g., molecules, proteins. In addition, we report graph classification performance on graphs with imputed features.
- **Type 2: Imputation on single-graph datasets.** We impute the missing values on single-graph datasets, e.g., social network. We report node classification performance on the graph with imputed features.

5.1 Imputation on Multi-Graph Datasets

Datasets. We conduct experiments on 6 benchmark datasets (Morris et al. 2020) from different domains: (1) bioinformatics, i.e., PROTEINS_full (Borgwardt et al. 2005) and ENZYMES (Schomburg et al. 2004); (2) chemistry, i.e., QM9 (Ramakrishnan et al. 2014) and FIRSTMM_DB (Neumann et al. 2013); (3) computer vision, i.e., FRANKENSTEIN (Orsini, Frasconi, and De Raedt 2015); (4) synthesis, i.e., Synthie (Morris et al. 2016). The detailed statistics is provided in Table 6 in Appendix E.1.

Baselines. We compare the performance of MEGAE against baselines in three categories: (1) statistical imputation methods including MEAN, KNN (Zhang 2012) and SVD (Troyanskaya et al. 2001); (2) deep learning-based imputation models including MICE (Van Buuren and Groothuis-Oudshoorn 2011), GAIN (Yoon, Jordon, and Schaar 2018), OT (Muzellec et al. 2020) and MIRACLE (Kyono et al. 2021); (3) graph learning-based models including GraphVAE (Simonovsky and Komodakis 2018), MolGAN (De Cao and Kipf 2018), GRAPE (You et al. 2020) and GDN (Li et al. 2021). For further details, please refer to Appendix E.3.

Setup. We use a 70-10-20 train-validation-test split and construct random missingness only on the test set. Each run has a different dataset split and the mask for feature missingness. After running for 5 trials, we report the Root Mean Squared Error (RMSE) results for imputation and mean accuracy on the test set for graph classification. For all baselines, we use a 2-layer GCN for downstream classification. For more setup details, please refer to Appendix E.2.

Results. In Table 1 we show the experiment results for 10% of missing features. Results indicate that our method has the lowest RMSE in the 5 real-world datasets and the synthetic

Methods	PubMed (RMSE)					AmaPhoto (RMSE)				
	0.1 Miss	0.3 Miss	0.5 Miss	0.7 Miss	0.99 Miss	0.1 Miss	0.3 Miss	0.5 Miss	0.7 Miss	0.99 Miss
sRMGCNN	0.0170	<u>0.0186</u>	0.0258	0.0309	0.0435	0.3191	0.3261	<u>0.3282</u>	<u>0.3359</u>	0.3660
GC-MC	0.0220	<u>0.0233</u>	0.0253	0.0270	0.0403	0.3154	0.3221	<u>0.3340</u>	<u>0.3574</u>	0.3913
GRAPE	0.0187	0.0198	<u>0.0234</u>	0.0303	0.0368	0.3207	0.3295	0.3357	0.3489	0.3896
VGAE	0.0196	0.0229	0.0298	0.0354	0.0525	0.3064	0.3196	0.3462	0.3684	0.3991
GDN	<u>0.0168</u>	0.0202	0.0239	<u>0.0254</u>	<u>0.0319</u>	<u>0.3006</u>	<u>0.3168</u>	0.3308	0.3372	<u>0.3617</u>
MEGAE	0.0157	0.0171	0.0185	0.0199	0.0218	0.2966	0.3121	0.3204	0.3243	0.3305
Performance gain	0.0011	0.0015	0.0049	0.0055	0.0101	0.0040	0.0047	0.0078	0.0116	0.0312
	 0.0063	 0.0062	 0.0113	 0.0155	 0.0307	 0.0241	 0.0174	 0.0258	 0.0441	 0.0686

Table 2: Mean RMSE results of attribute imputation with different missing rates on PubMed and AmaPhoto.

Methods	PubMed (Acc. = 77.11% with full features)					AmaPhoto (Acc. = 92.03% with full features)				
	0.1 Miss	0.3 Miss	0.5 Miss	0.7 Miss	0.99 Miss	0.1 Miss	0.3 Miss	0.5 Miss	0.7 Miss	0.99 Miss
GCNMF	75.96%	74.42%	74.03%	70.09%	41.46%	91.64%	91.46%	91.05%	<u>90.16%</u>	<u>73.03%</u>
PaGNN	<u>76.55%</u>	<u>75.77%</u>	<u>74.45%</u>	72.15%	<u>43.63%</u>	91.81%	<u>91.70%</u>	91.02%	90.12%	71.51%
sRMGCNN	75.96%	74.74%	72.79%	71.14%	21.17%	91.36%	91.19%	91.02%	90.12%	66.56%
GC-MC	74.44%	74.06%	72.87%	72.67%	36.69%	91.46%	91.26%	90.22%	89.82%	60.13%
GRAPE	75.53%	74.51%	73.65%	71.75%	38.18%	91.22%	91.05%	90.33%	89.63%	68.28%
VGAE	75.21%	73.69%	71.64%	69.96%	16.66%	91.53%	91.43%	89.23%	88.82%	63.24%
GDN	75.87%	74.22%	73.36%	<u>73.15%</u>	38.93%	<u>91.83%</u>	91.63%	<u>91.10%</u>	90.01%	66.62%
MEGAE	76.66%	76.30%	75.03%	74.47%	49.58%	91.92%	91.85%	91.41%	90.58%	81.88%
Performance gain	0.11%	0.53%	0.58%	1.32%	5.95%	0.09%	0.15%	0.31%	0.42%	8.85%
	 1.45%	 2.61%	 3.39%	 4.51%	 32.92%	 0.70%	 0.80%	 2.18%	 1.76%	 21.75%

Table 3: Test accuracy with different missing rates on PubMed and AmaPhoto node classification benchmarks.

dataset. Moreover, almost all graph learning-based methods outperform non-graph learning ones, demonstrating the effectiveness of modeling the dependencies between graph nodes. In addition, we introduce two more baselines – Inv_MLP denotes our tight wavelet encoder with Multi-Layer Perceptron (MLP) decoder (Simonovsky and Komodakis 2018) and Inv_GCN denotes our encoder with GCN decoder (Kipf and Welling 2017). MEGAE achieves 12.4% lower RMSE than the second best baselines on average over all the datasets. For a more comprehensive and realistic consideration for missing mechanisms, we provide the evaluation for evaluating MEGAE in two less common scenarios. We obtain similar gains across three missing mechanisms, MEGAE obtains the lowest errors in all cases, which shows the strong generalization among the three missing mechanisms.

As for downstream classifications, we show the test accuracy of different models on imputed graph features with 10% missing. We observe that smaller imputation errors are more likely to lead to higher test accuracy. This is relatively expected, as data with less distribution shift easily yields better representation. MEGAE outperforms all other baselines on accuracy conditioning on the effectiveness imputation results.

5.2 Imputation on Single-Graph Datasets

Dataset. We evaluate the effectiveness of MEGAE on four commonly used datasets including three citation network datasets—Cora, Citeseer, Pubmed (Sen et al. 2008) and a co-purchase network dataset—AmaPhoto (McAuley et al. 2015). We show the dataset statistics in Table 7 in Appendix E.1.

Baselines. We compare the imputation performance of MEGAE against competitive methods including sRMGCNN (Monti, Bronstein, and Bresson 2017), GC-MC (Berg, Kipf, and Welling 2017), GRAPE (You et al. 2020), VGAE (Kipf and Welling 2016) and GDN (Li et al. 2021). For evaluating node classification performance of MEGAE, we additionally experiment with GCNMF (Taguchi, Liu, and Murata 2021) and PaGNN (Jiang and Zhang 2020) that directly implement graph learning without imputation.

Setup. We closely follow (Kipf and Welling 2017) to perform standard dataset split. For all datasets, each run has a different train-validation-test split and the mask for random missingness across all feature dimensions. After running for 5 trails, we report the results of mean RMSE for imputation and mean accuracy on the test set for node classification. For all baselines, we use a 2-layer GCN for downstream

classification.

Results. Table 2 shows the mean RMSE results for imputation under different missing rates on PubMed and AmaPhoto, we also provide the results on Cora and Citeseer in Table 8 in Appendix. MEGAE consistently outperforms other baselines. Moreover, MEGAE tends to have greater performance gains when faced with a more severe missingness, which demonstrates MEGAE’s robustness. It is noteworthy that under an extremely high missing rate (0.99), MEGAE obtains margins of 31.66% on PubMed, 8.63% on AmaPhoto, 34.24% on Cora and 19.54% on Citeseer compared to the second best imputation methods. In Table 3, we show the mean accuracy for node classification on PubMed and AmaPhoto with imputed attributes (results on Cora and Citeseer are in Table 9 in Appendix). We observe that most methods work well with less than 0.3 missing rates. Moreover, MEGAE’s robustness and effectiveness are better as its accuracy gains tend to be positively correlated with missing rates. With the consistent imputation improvements on all datasets, MEGAE has accuracy gains of 5.95% on PubMed, 8.85% on AmaPhoto, 7.90% on Cora and 6.97% on Citeseer under a 0.99 missing rate.

5.3 Effectiveness of Wavelet and Regularization

Graph Wavelet Autoencoders. We study the role of tight wavelet frames with the configuration of regularization in MEGAE. Another two paradigms are deployed to get wavelet frames $\hat{g} = \{\hat{g}_1, \dots, \hat{g}_M\}$, including GWNN (Xu et al. 2019) with heat kernels and SGWT (Hammond, Vandergheynst, and Gribonval 2011) with highly-localized kernels.

Dataset	GWNN	SGWT	Ours
ENZYMES	+0.0042	+0.0044	0.0223
PROTEINS_full	+0.0070	+0.0052	0.0099
QM9	+0.0102	+0.0166	0.1396
Synthetic	+0.0203	+0.0126	0.1203
FRANKENSTEIN	+0.0103	+0.0185	0.0936
FIRST_DB	+0.0036	+0.0073	0.0789

Table 4: RMSE with three graph wavelet paradigms. We report the mean RMSE of MEGAE and show the increase w.r.t. RMSE of GWNN and SGWT compared to ours.

Here we provide experiments on 6 multi-graph datasets under 0.1 missing rate (in Table 4), reporting the increases w.r.t. RMSE using another two graph wavelet paradigms compared to ours, all three wavelet paradigms have 9 channels. Also, we execute experiments on 2 datasets (in Figure 3) assembling various channel numbers on three graph wavelet paradigms. Three paradigms all share the same computational complexity $O(K \times |\mathbb{E}|)$, where K is the order of polynomials and $|\mathbb{E}|$ is the number of edges. Table 4 verifies the effectiveness of the proposed tight wavelet kernels. Figure 3 shows the influence of the number of wavelet kernels M . We observe that a larger M value harms the performance of GAEs with GWNN and SGWT paradigms. As for our proposed MEGAE, the RMSE first drops and then increases along w.r.t. the number of wavelet kernels, which shows there is a tradeoff between entropy approximation error and model generalization.




Figure 3: RMSE with different numbers of wavelet kernels. We show how the kernel number influences the performance of GAEs with different wavelet paradigms.

Entropy Regularization. We study the role of entropy regularization with the configuration of tight wavelet in MEGAE. We execute experiment on 6 datasets under 10% missing rates. In Table 5 we provide the increase of RMSE by removing entropy regularization and the change in graph spectral entropy w.r.t. the original graph spectral entropy (denoted as Δ). We observe that GAE without regularization will seriously lead to a decrease in spectral entropy (i.e., aforementioned spectral concentration) of the imputed graphs, and more importantly, \mathcal{L}_S improves the performance of the proposed MEGAE with a substantially maintenance of the real entropy.

Dataset	\mathcal{L}_R	Δ_1	$\mathcal{L}_R + \mathcal{L}_S$	Δ_2
ENZYMES	+0.0069	-29%	0.0223	+3%
PROTEINS_full	+0.0061	-26%	0.0099	+2%
QM9	+0.0188	-18%	0.1396	-1%
Synthetic	+0.0316	-23%	0.1203	+3%
FRANKENSTEIN	+0.0229	-33%	0.0936	-2%
FIRSTMM_DB	+0.0086	-21%	0.0789	-2%

Table 5: RMSE with/without entropy regularization. Δ_1 and Δ_2 denote the entropy changes with/without the regularization w.r.t. the original graph spectral entropy.

Discussion. The effectiveness of our proposed tight wavelet and entropy regularization has been shown in Table 4 and Table 5, from which we can easily find that both have obvious improvements for decreasing imputation errors. Importantly, we believe the entropy regularization contributes the most to the performance of MEGAE. By analogizing MEGAE to the variational autoencoders, we provide intuitions on why MEGAE works. The proof can be found in Appendix D.

6 Conclusion.

In this paper we show that graph spectral entropy is a good optimizable proxy for graph attribute imputation. Motivated by eliminating the spectral concentration of GAEs and retaining all spectral components, we propose MEGAE with a graph spectral entropy regularization for imputation. We develop a theoretical guarantee that our proposed MEGAE could maximize the target entropy without eigen-decomposition. Experimental results of various datasets on imputation and downstream classification tasks show the effectiveness of MEGAE for handling missing attributes on graphs.

Acknowledgements

The research of Li was supported by NSFC Grant No. 62206067, Tencent AI Lab Rhino-Bird Focused Research Program RBFR2022008 and Guangzhou-HKUST(GZ) Joint Funding Scheme. The research of Tsung was supported by the Hong Kong RGC General Research Funds 16216119 and Foshan HKUST Projects FSUST20-FYTRI03B.

References

- Ahmad, O.; and Sheikh, N. A. 2018. On characterization of nonuniform tight wavelet frames on local fields. *Anal. Theory Appl*, 34(2): 135–146.
- Berg, R. v. d.; Kipf, T. N.; and Welling, M. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*.
- Borgwardt, K. M.; Ong, C. S.; Schönaauer, S.; Vishwanathan, S.; Smola, A. J.; and Krieger, H.-P. 2005. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1): i47–i56.
- Burch, S. F.; Gull, S. F.; and Skilling, J. 1983. Image restoration by a powerful maximum entropy method. *Computer Vision, Graphics, and Image Processing*, 23(2): 113–128.
- Cai, D.; He, X.; and Han, J. 2007a. Spectral regression for efficient regularized subspace learning. In *2007 IEEE 11th international conference on computer vision*, 1–8. IEEE.
- Cai, D.; He, X.; and Han, J. 2007b. SRDA: An efficient algorithm for large-scale discriminant analysis. *IEEE transactions on knowledge and data engineering*, 20(1): 1–12.
- Candès, E. J.; and Recht, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6): 717–772.
- De Cao, N.; and Kipf, T. 2018. MolGAN: An implicit generative model for small molecular graphs. In *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- Donnat, C.; Zitnik, M.; Hallac, D.; and Leskovec, J. 2018. Learning structural node embeddings via diffusion wavelets. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1320–1329.
- Frieden, B. R. 1972. Restoring with maximum likelihood and maximum entropy. *JOSA*, 62(4): 511–518.
- Hammond, D. K.; Vandergheynst, P.; and Gribonval, R. 2011. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2): 129–150.
- Hastie, T.; Mazumder, R.; Lee, J. D.; and Zadeh, R. 2015. Matrix completion and low-rank SVD via fast alternating least squares. *The Journal of Machine Learning Research*, 16(1): 3367–3402.
- Jia, H.; Ding, S.; Du, M.; and Xue, Y. 2016. Approximate normalized cuts without Eigen-decomposition. *Information Sciences*, 374: 135–150.
- Jiang, B.; and Zhang, Z. 2020. Incomplete graph representation and learning via partial graph neural networks. *arXiv preprint arXiv:2003.10130*.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *NeurIPS-Workshop*.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. *ICLR*.
- Kyono, T.; Zhang, Y.; Bellot, A.; and van der Schaar, M. 2021. MIRACLE: Causally-Aware Imputation via Learning Missing Data Mechanisms. In *NeurIPS*.
- Leonardi, N.; and Van De Ville, D. 2013. Tight wavelet frames on multislice graphs. *IEEE Transactions on Signal Processing*, 61(13): 3357–3367.
- Li, J.; Li, J.; Liu, Y.; Yu, J.; Li, Y.; and Cheng, H. 2021. Deconvolutional Networks on Graph Data. In *NeurIPS*.
- Li, J.; Yu, T.; Juan, D.-C.; Gopalan, A.; Cheng, H.; and Tomkins, A. 2020. Graph Autoencoders with Deconvolutional Networks. *arXiv preprint arXiv:2012.11898*.
- Li, M.; Guo, X.; Wang, Y.; Wang, Y.; and Lin, Z. 2022. G²CN: Graph Gaussian Convolution Networks with Concentrated Graph Filters. In *International Conference on Machine Learning*, 12782–12796. PMLR.
- McAuley, J.; Targett, C.; Shi, Q.; and Van Den Hengel, A. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 43–52.
- Monti, F.; Boscaini, D.; Masci, J.; Rodola, E.; Svoboda, J.; and Bronstein, M. M. 2017. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5115–5124.
- Monti, F.; Bronstein, M.; and Bresson, X. 2017. Geometric matrix completion with recurrent multi-graph neural networks. *Advances in neural information processing systems*, 30.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. Tudataset: A collection of benchmark datasets for learning with graphs. In *ICML Workshop on Graph Representation Learning and Beyond*.
- Morris, C.; Kriege, N. M.; Kersting, K.; and Mutzel, P. 2016. Faster kernels for graphs with continuous attributes via hashing. In *ICDM*, 1095–1100. IEEE.
- Murray, J. S.; and Reiter, J. P. 2016. Multiple imputation of missing categorical and continuous values via Bayesian mixture models with local dependence. *Journal of the American Statistical Association*, 111(516): 1466–1479.
- Muzellec, B.; Josse, J.; Boyer, C.; and Cuturi, M. 2020. Missing data imputation using optimal transport. In *ICML*, 7130–7140. PMLR.
- Narayan, R.; and Nityananda, R. 1986. Maximum entropy image restoration in astronomy. *Annual review of astronomy and astrophysics*, 24(1): 127–170.
- Neumann, M.; Moreno, P.; Antanas, L.; Garnett, R.; and Kersting, K. 2013. Graph kernels for object category prediction in task-dependent robot grasping. In *MLG*, 0–6.
- Nt, H.; and Maehara, T. 2019. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*.
- Orsini, F.; Frasconi, P.; and De Raedt, L. 2015. Graph invariant kernels. In *IJCAI*.

- Park, J.; Lee, M.; Chang, H. J.; Lee, K.; and Choi, J. Y. 2019. Symmetric graph convolutional autoencoder for unsupervised graph representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6519–6528.
- Ramakrishnan, R.; Dral, P. O.; Rupp, M.; and Von Lilienfeld, O. A. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1): 1–7.
- Roth, P. L. 1994. Missing data: A conceptual review for applied psychologists. *Personnel psychology*, 47(3): 537–560.
- Salomon, A.; Cahen, D.; Lindsay, S.; Tomfohr, J.; Engelkes, V. B.; and Frisbie, C. D. 2003. Comparison of electronic transport measurements on organic molecules. *Advanced Materials*, 15(22): 1881–1890.
- Schomburg, I.; Chang, A.; Ebeling, C.; Gremse, M.; Heldt, C.; Huhn, G.; and Schomburg, D. 2004. BRENDA, the enzyme database: updates and major new developments. *Nucleic acids research*, 32(suppl_1): D431–D433.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI magazine*, 29(3): 93–93.
- Shuman, D. I.; Wiesmeyer, C.; Holighaus, N.; and Vandergheynst, P. 2015. Spectrum-adapted tight graph wavelet and vertex-frequency frames. *IEEE Transactions on Signal Processing*, 63(16): 4223–4235.
- Simonovsky, M.; and Komodakis, N. 2018. GraphVAE: Towards generation of small graphs using variational autoencoders. In *ICANN*, 412–422. Springer.
- Sterne, J. A.; White, I. R.; Carlin, J. B.; Spratt, M.; Royston, P.; Kenward, M. G.; Wood, A. M.; and Carpenter, J. R. 2009. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. *BMJ*, 338.
- Taguchi, H.; Liu, X.; and Murata, T. 2021. Graph convolutional networks for graphs containing missing features. *Future Generation Computer Systems*, 117: 155–168.
- Tang, J.; Li, J.; Gao, Z.; and Li, J. 2022. Rethinking Graph Neural Networks for Anomaly Detection. In *International Conference on Machine Learning*.
- Troyanskaya, O.; Cantor, M.; Sherlock, G.; Brown, P.; Hastie, T.; Tibshirani, R.; Botstein, D.; and Altman, R. B. 2001. Missing value estimation methods for DNA microarrays. *Bioinformatics*, 17(6): 520–525.
- Tu, W.; Zhou, S.; Liu, Y.; and Liu, X. 2021. Siamese Attribute-missing Graph Auto-encoder. *arXiv preprint arXiv:2112.04842*.
- Van Buuren, S. 2018. *Flexible imputation of missing data*. CRC press.
- Van Buuren, S.; and Groothuis-Oudshoorn, K. 2011. mice: Multivariate imputation by chained equations in R. *Journal of statistical software*, 45: 1–67.
- White, I. R.; Royston, P.; and Wood, A. M. 2011. Multiple imputation using chained equations: issues and guidance for practice. *Statistics in medicine*, 30(4): 377–399.
- Wooldridge, J. M. 2007. Inverse probability weighted estimation for general missing data problems. *Journal of econometrics*, 141(2): 1281–1301.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*, 6861–6871. PMLR.
- Xia, J.; Zhang, S.; Cai, G.; Li, L.; Pan, Q.; Yan, J.; and Ning, G. 2017. Adjusted weight voting algorithm for random forests in handling missing values. *Pattern Recognition*, 69: 52–60.
- Xu, B.; Shen, H.; Cao, Q.; Cen, K.; and Cheng, X. 2020. Graph convolutional networks using heat kernel for semi-supervised learning. *arXiv preprint arXiv:2007.16002*.
- Xu, B.; Shen, H.; Cao, Q.; Qiu, Y.; and Cheng, X. 2019. Graph Wavelet Neural Network. In *ICLR*.
- Yomogida, Y.; Pu, J.; Shimotani, H.; Ono, S.; Hotta, S.; Iwasa, Y.; and Takenobu, T. 2012. Ambipolar Organic Single-Crystal Transistors Based on Ion Gels. *Advanced Materials*, 24(32): 4392–4397.
- Yoon, J.; Davtyan, C.; and van der Schaar, M. 2016. Discovery and clinical decision support for personalized healthcare. *IEEE Journal of Biomedical and Health Informatics*, 21(4): 1133–1145.
- Yoon, J.; Jordon, J.; and Schaar, M. 2018. Gain: Missing data imputation using generative adversarial nets. In *ICML*, 5689–5698. PMLR.
- Yoon, S.; and Sull, S. 2020. GAMIN: generative adversarial multiple imputation network for highly missing data. In *CVPR*, 8453–8461.
- You, J.; Ma, X.; Ding, D. Y.; Kochenderfer, M.; and Leskovec, J. 2020. Handling missing data with graph representation learning. In *NeurIPS*.
- Zhang, M.; and Chen, Y. 2019. Inductive matrix completion based on graph neural networks. *arXiv preprint arXiv:1904.12058*.
- Zhang, S. 2012. Nearest neighbor selection for iteratively kNN imputation. *Journal of Systems and Software*, 85(11): 2541–2552.