Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

Embedding graphs on Grassmann manifold

Bingxin Zhou^{a,b,*,1}, Xuebin Zheng^{a,1}, Yu Guang Wang^{b,c}, Ming Li^d, Junbin Gao^a

^a Discipline of Business Analytics, The University of Sydney Business School, The University of Sydney, NSW 2006, Australia

^b Institute of Natural Sciences and School of Mathematical Sciences, Shanghai Jiao Tong University, Shanghai 200240, China

^c School of Mathematics and Statistics, The University of New South Wales, NSW 2052, Australia

^d Key Laboratory of Intelligent Education Technology and Application of Zhejiang Province, Zhejiang Normal University, Jinhua 321004, China

ARTICLE INFO

ABSTRACT

Article history: Received 4 November 2021 Received in revised form 25 March 2022 Accepted 2 May 2022 Available online 10 May 2022

Keywords: Grassmann manifold Graph neural network Projection embedding Subspace clustering Learning efficient graph representation is the key to favorably addressing downstream tasks on graphs, such as node or graph property prediction. Given the non-Euclidean structural property of graphs, preserving the original graph data's similarity relationship in the embedded space needs specific tools and a similarity metric. This paper develops a new graph representation learning scheme, namely EGG, which embeds approximated second-order graph characteristics into a Grassmann manifold. The proposed strategy leverages graph convolutions to learn hidden representations of the corresponding subspace of the graph, which is then mapped to a Grassmann point of a low dimensional manifold through truncated singular value decomposition (SVD). The established graph embedding approximates denoised correlationship of node attributes, as implemented in the form of a symmetric matrix space for Euclidean calculation. The effectiveness of EGG is demonstrated using both clustering and classification tasks at the node level and graph level. It outperforms baseline models on various benchmarks.

© 2022 Elsevier Ltd. All rights reserved.

1. Introduction

As one of the most prominent avenues in geometric deep learning, graph neural networks (GNNs) have received growing attention over the last few years (Bronstein, Bruna, LeCun, Szlam, & Vandergheynst, 2017; Hamilton, 2020; Wu et al., 2020; Zhang, Cui, & Zhu, 2020; Zhou et al., 2020). Common to many GNN-based predictive tasks, distilling key features and structural information from the given graph data stays within the core of designing an effective graph representation learning.

Graph convolution (Bruna, Zaremba, Szlam, & Lecun, 2014), especially graph neural message passing (Gilmer, Schoenholz, Riley, Vinyals, & Dahl, 2017), provides an efficient expression to the information flow of the underlying graph through aggregating regional features over the node neighborhood. For a set of multiple graphs of varying sizes and topological structure, employing arbitrary graph convolutions fails to coincide with the size of graph representation. Instead, an appropriate graph pooling scheme is required to establish graph-level representations of a uniform scale. Furthermore, a handful of pooling strategies have been proposed to scale down the graph embedding by extracting the key components of the graph representation. Depending

* Corresponding author at: Discipline of Business Analytics, The University of Sydney Business School, The University of Sydney, NSW 2006, Australia.

E-mail address: bzho3923@uni.sydney.edu.au (B. Zhou).

¹ Both authors contribute equally.

https://doi.org/10.1016/j.neunet.2022.05.001 0893-6080/© 2022 Elsevier Ltd. All rights reserved. on whether the hidden attributes are coarsened along with the adjacency matrix, different strategies are categorized as either global (Wang & Ji, 2020; Xu, Hu, Leskovec, & Jegelka, 2019; Zheng et al., 2021) or hierarchical pooling (Gao & Ji, 2019; Ma, Wang, Aggarwal, & Tang, 2019; Zheng, Zhou, Li, Wang, & Gao, 2020).

While each operation designs its unique standard for graph coarsening and feature extraction, common to all pooling strategies is the requirement of node permutation invariance. When employing feature extraction, the node order of an undirected graph must not incline the network to excessively concentrate on individual attributes. The permutation invariance property is intuitive in heuristic pooling operations like summation, averaging, and maximization, where the aggregation rules of regional patterns do not rely on the arrangement of nodes sequence. One possibility to define a permutation invariant pooling operation on graph topology is to view nodes of a graph as a set of elements. For instance, Kolouri, Naderializadeh, Rohde, and Hoffmann (2021) establish the Wasserstein embedding of a node set under the linear optimal transport framework (Mémoli, 2011). A proper Wasserstein metric is designed to match pairs of node sets while avoiding a direct node selection criterion, as the hierarchical pooling schemes.

This work instead studies the expression of the node set of a graph through *manifold learning* (Absil, Mahony, & Sepulchre, 2008). Recall that the fundamental assumption of a smooth graph in the design of graph convolutions is that spatially connected nodes are likely to share similar characteristics and the nodes









Fig. 1. Computational principle of the proposed framework EGG. (a) Given a graph \mathcal{G}_i , i.e., a set of nodes, the target is to train a model that assigns a label to each of them. (b) Every \mathcal{G}_i rectifies a subspace of its most representative bit on the manifold space, where their geodesic distance is measured by their principal angle. (c) From the perspective of an orthonormal basis, these subspaces can be embedded as Grassmann points of a Grassmann manifold, where similar points have a small distance. (d) These Grassmann points support an easy projection operation to the space of symmetric matrices for deep learning tasks, such as classification and clustering.

from the same class are likely to have similar attributes. Naturally, the hidden representation of an attributed graph creates a subspace of lower dimension, or equivalently, a point of a Grassmann manifold. Consequently, a sophisticated learning task over graphs of varying size and topology accomplishes its transformation to a new learning task over Grassmann points of a fixed-dimensional Grassmann manifold.

We name the above embedding strategy of mapping graphs to Grassmann points as **E**mbedding **G**raphs on a **G**rassmann manifold, or EGG for short. This framework is a generic method to express graphs with a Grassmann manifold subspace analysis. As demonstrated in Fig. 1, each representation of node sets establishes a subspace of indeterminate dimensions. EGG then embeds these subspaces to Grassmann points of a Grassmannian \mathcal{M} , where every point is explicitly represented by an orthonormal matrix. Furthermore, these Grassmann points support an effortless inversion to the Euclidean space through symmetric projection, where new representations are of the same dimension, and they are ready for conventional graph classification or segmentation tasks.

In comparison to existing graph distilling strategies, this new architecture design is one practice of exploiting principal components of graphs from non-linear information transformation. Each node community formulates a subspace of coarsened and smoother higher-level expression. While conventional graph aggregation generally requires stacked convolution or pooling layers to allow non-linear propagation, EGG leverages the truncated singular value decomposition (trSVD) to directly compress the principal components and construct the smooth subspace. In addition, the projected results from these acquired graph Grassmann embeddings approximate the second-order covariance of node attributes, which gains more expressive power than typical first-order expression from the conventional aggregation and distilling operations. Moreover, the proposed EGG for graph-level tasks guarantees the critical property of permutation invariance, which is an essential requirement of a qualified graph pooling design, yet it has been ignored by many existing methods.

The preliminary idea of the developed framework EGG was first introduced in a workshop paper (Zhou, Zheng, Wang, Li, & Gao, 2021). This extension provides abundant details for understanding the rationale and paradigm of the proposed graph embedding scheme. Furthermore, the embedding strategy is expanded from graph pooling applications to more general scenarios, where in this complete work we exploit the possibility of handling lower-level unsupervised learning tasks of node segmentation. Additional investigations are addressed to interpret the learned expression and avoid the black-box model design.

The rest of this paper is organized as follows. Section 2 reviews the previous literature on graph representation learning and Grassmann deep learning applications. Section 3 introduces the Grassmann geometry that is closely related. Section 4 details the two critical ingredients of analyzing subspaces in a

Grassmannian. We then demonstrate our methods with two specific applications: graph classification and node clustering. The problems are formulated in Section 5, and the empirical performances are reported in Section 6. Further investigations on the significance of EGG are addressed in Section 7.

2. Related work

An attributed undirected graph is denoted as $\mathcal{G}_i = (V_i, E_i, X_i)$ of $n_i := |V_i|$ nodes and $|E_i|$ edges. The node is featured by $X_i \in \mathbb{R}^{n_i \times d}$, and the (weighted) edges for the structure information are described by an adjacency matrix $A \in \mathbb{R}^{n_i \times n_i}$. Since the graph topology provides additional information, graph representation learning aims at encoding such a structural expression to conventional vector representations for deep learning models that assign labels to instances. A node-level graph learning task assigns a label $\{y_i\}$ to each node of the graph \mathcal{G}_i , while a graphlevel task finds a sequence of N labels $\{y_1, y_2, \ldots, y_N\}$ from a set of input graphs $\mathbb{G} = \{\mathcal{G}_1, \ldots, \mathcal{G}_N\}$. Depending on the nature of the assigned labels, the learning task can be categorized to either regression or classification.

Spatial graph convolution. The emerging development of graph neural networks (GNNs) generates enormous work for graph representation learning. Typically, the topological embedding is realized by graph convolutional layers. A spatial-based propagation rule (Gilmer et al., 2017) leverages proper feature extraction and aggregation from the central node's local community or the neighborhood. The propagation rule can be designed as flexible as weighted average (Bo, Wang, Shi, & Shen, 2021; He et al., 2020; Kipf & Welling, 2017), concatenation (Hamilton, Ying, & Leskovec, 2017; Xu et al., 2018), learnable attention (Kim & Oh, 2021; Veličković et al., 2018; Wang et al., 2019), or other adaptive choices (Brockschmidt, 2020; Tailor, Opolka, Lio, & Lane, 2021). To allow a broader receptive field, multiple convolution layers are frequently stacked for multi-hop neighborhood aggregation. Nevertheless, the majority of the aggregation rules are carried out in the first-order space, which omits the second-order covariance information that can capture insightful non-linear relationships of the feature attributes (Lin, RoyChowdhury, & Maji, 2015; Wang, Long, Wang, & Yu, 2017).

Graph pooling and down-sampling. As a graph-level learning task involves multiple graphs of a diverse number of nodes, it is crucial for GNNs to unify the dimension of the output graph representation by pooling operations. For instance, TopKPooL (Cătălina, Veličković, Jovanović, Kipf, & Liò, 2018) formulates a score function to rank graph nodes and picks the parent nodes of subgraphs from graph clusters to hierarchically coarsen a graph. Other research enhances the selection efficacy through a carefully designed scoring mechanism, such as multilayer perceptron (MLP) (Lee, Lee, & Kang, 2019) or attention (Knyazev, Taylor, & Amer, 2019; Li, Tarlow, Brockschmidt, & Zemel, 2016). The authors of Gao and Ji (2019), Ma et al. (2019), Zheng et al. (2020) extended graph coarsening conventions to different slicing principles, although the role of graph clustering and its influence on local pooling has been challenged by the literature (Mesquita, Souza, & Kaski, 2020; Wang & Ji, 2020). Alternatively, global graph pooling strategies are pumped out in practice with a simpler design and comparable performance (Xu et al., 2019; Zheng et al., 2021).

Grassmann manifold in deep learning. Grassmann manifolds play an important role in recommender systems (Boumal & Absil, 2015; Dai, Kerman, & Milenkovic, 2012), computer vision (Lui, 2012; Minh, Murino, & Minh, 2016) and pattern recognition (Huang, Wang, Shan, & Chen, 2015; Slama, Wannous, Daoudi, & Srivastava, 2015). Grassmann learning exploits subspace-invariant features and harnesses the structural information of sample sets, which improves the prediction performance of a model with lower complexity and higher robustness. Due to these privileges, the Grassmann manifold is often approached as a tool of nonlinear dimensionality reduction (Dong, Frossard, Vandergheynst, & Nefedov, 2014; Koch & Lubich, 2007; Ngo & Saad, 2012; Zhou, Gao, Tran, & Gerlach, 2021) or optimization objectives (Absil, Mahony, & Sepulchre, 2004; Edelman, Arias, & Smith, 1998). While direct computations on a Grassmannian can be sophisticated, other research investigates pipeline Grassmann points to a Grassmann learning algorithm, such as Deep Grassmann Networks (Huang, Wu, & Van Gool, 2018) and Grassmann clustering (Wang, Hu, et al., 2017; Wang, Hu, Gao, Sun, & Yin, 2014).

3. Grassmann geometry

This section overviews the mathematical formulation of the Grassmann points and Grassmann manifold. We also discuss the measurement on the geodesic distance for Grassmann points, as well as their Euclidean counterparts.

3.1. Grassmann manifold

Grassmann manifold is a manifold of matrices of a specific rank. Each Grassmann point is represented by an orthonormal matrix, and it corresponds to a subspace of the underlying real Euclidean space. To say it precisely,

Definition 1 (*Grassmann Manifold (Absil et al., 2008*)). The Grassmann manifold $\mathcal{M}(p, m)$ ($p \leq m$) consists of all *p*-dimensional subspaces of the Euclidean space \mathbb{R}^m , i.e.,

$\mathcal{M}(p, m) = \{\mathcal{U} \subset \mathbb{R}^m : \mathcal{U} \text{ is a subspace, } \dim(\mathcal{U}) = p\}.$

A particular Grassmann point $\mathcal{U} \in \mathcal{M}(p, m)$ is identified by an orthonormal matrix $\mathbf{U} \in \mathbb{R}^{m \times p}$, which is an equivalence class of all rank-*p* matrices that spans \mathcal{U} . That is, $\mathcal{M}(p, m) = \{\text{span}(\mathbf{U}) : \mathbf{U} \in \mathbb{R}^{m \times p}, \mathbf{U}^{\top}\mathbf{U} = \mathbb{I}_p\}$. Furthermore, this subspace is identified uniquely by a projector $\Pi(\mathbf{U})$ on \mathcal{U} , such that $\Pi(\mathbf{U}) = \mathbf{U}\mathbf{U}^{\top}$. The Grassmann manifold is an abstract quotient manifold that one can represent in many ways, such as the Lie group theory (Bendokat, Zimmermann, & Absil, 2020; Gallivan, Srivastava, Liu, & Van Dooren, 2003). To best allow convenient algebraic calculations, this work constructs Grassmann points from the perspective of projection matrices (Chikuse, 2003).

3.2. Subspace distance

The distance between two Grassmann points is measured differently from the conventional Euclidean metric due to the curvature of the Grassmannian. The *geodesic distance* is thus defined as the length of the shortest path along the manifold between two points, which is a function of the principal angles of the two subspaces or analogously the two Grassmann points, as we introduce now.

Definition 2 (*Principal Angle*). Given two Grassmann points $U_1, U_2 \in \mathcal{M}(p, m)$ and their orthonormal bases $\mathbf{U}_1 = [(\mathbf{u}_1)_1, \dots, (\mathbf{u}_1)_p]$, $\mathbf{U}_2 = [(\mathbf{u}_2)_1, \dots, (\mathbf{u}_2)_p] \in \mathbb{R}^{m \times p}$, we define their principal angles $0 \le \theta_1 \le \dots \le \theta_p \le \frac{\pi}{2}$ recursively by

 $\cos(\theta_i) = \max(\boldsymbol{u}_1)_i^{\top}(\boldsymbol{u}_2)_i$

s.t. $\|\boldsymbol{u}_1\|_2 = \|\boldsymbol{u}_2\|_2 = 1$, $(\boldsymbol{u}_1)_i^\top (\boldsymbol{u}_1)_j = (\boldsymbol{u}_2)_i^\top (\boldsymbol{u}_2)_j = 0 \quad \forall j < i$.

The principal angles describe the smallest p angles between all possible bases of the two p-dimensional subspaces (U_1 and U_2). With a sequence of principal angles $\Theta = [\theta_1, \ldots, \theta_p]$, the geodesic distance between the two Grassmann points is a function of the principal angles, i.e., $d(U_1, U_2) = ||\Theta||_2$.

In literature, there exist many other measurements to describe the discrepancy between subspaces, so that a closed-form solution for optimization on a Grassmann manifold becomes possible. For example, the *projection distance* (Chikuse, 2003) embeds a Grassmann manifold $\mathcal{M}(p, m)$ into a higher *m* dimensional Euclidean space in the form of a Symmetric Positive-Definite (SPD) matrix; the chordal distance and the Procrustes distance (Ye & Lim, 2016) measures the total squared sine angle between Grassmann points, which is usually used for shape analysis. This work follows the first measure of the projection distance, which also supports kernelized Grassmann learning (Harandi, Salzmann, Jayasumana, Hartley, & Li, 2014) and has been well-explored in learning low-rank approximation (Dong et al., 2014; Ngo & Saad, 2012; Zhou, Gao, et al., 2021) and pattern recognition (Huang et al., 2015; Slama et al., 2015).

4. Grassmannian subspace analysis on graphs

This section provides a guideline for graph smoothing and distilling through node sets embedding to a continuous and smooth Grassmann manifold. As we shall introduce below, our proposed method rectifies graph representations with embedded structure information to a Grassmannian of their feature space at a lower dimension. The Euclidean representations of these Grassmann instances from the projection perspective can be considered as an approximate version of the feature correlations that eliminates unnecessary variances. A Grassmann embedding appends non-linear smoothing effects to the graph representations that are usually achieved by stacking up fully-connected layers. The output representation meets the key requirements of a standard graph embedding scheme, and it allows arbitrary computations adapted to a Euclidean space.

4.1. Problem formulation

We begin with a set of hidden representations $\mathbb{H} = \{H_1, \ldots, H_N\}$ from graph convolutional layers for a given set of *N* graphs $\mathbb{G} = \{\mathcal{G}_1, \ldots, \mathcal{G}_N\}$. Here $H_i \in \mathbb{R}^{n_i \times m}$ is with respect to n_i nodes in \mathcal{G}_i and a number of *m* hidden neurons for the last graph convolution operation. The H_i for graph-level learning tasks requires a graph representation that is irrelevant to the node size and has an identical dimension to other $H \in \mathbb{H}$. Therefore, at the rectification step, each graph representation of \mathbb{H} to a Grassmann point is identified with an orthonormal basis of H, and is aligned to the same Grassmannian.

4.2. Manifold rectification

Suppose we have obtained a graph hidden representation $\mathbf{H} \in \mathbb{H} \subset \mathbb{R}^{n \times m}$ as well as its row-generated subspace span(\mathbf{H}^{\top}), which as mentioned can be achieved by employing one or multiple layers of graph convolution. Our target is to find a concrete Grassmann representation with most variations of the data. We characterize the representation as an orthogonal basis of the subspace, which can be rectified in several ways, such as the QR decomposition of matrix \mathbf{H}^{\top} , as demonstrated by Huang et al. (2018). Here we consider another classic method of the truncated singular value decomposition (SVD) to find the best low-rank approximation of the hidden feature space \mathbf{H}^{\top} in the sense of the least-squares (Chikuse, 2003).

The preliminary goal of employing the manifold embedding is to establish a graph representation of a unit dimension k. We therefore leverage the truncated SVD on \boldsymbol{H} to obtain the most representative basis $\boldsymbol{U} = [\boldsymbol{u}_1, \boldsymbol{u}_2, \dots, \boldsymbol{u}_p]$ of the subspace $\operatorname{span}(\boldsymbol{H}^{\top})$, i.e.,

$$\boldsymbol{H}^{\top} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\top}, \tag{1}$$

where the Grassmann point $\mathcal{U} = [\mathbf{U}]$ is an equivalence class of \mathbf{U} . The $\mathbf{U} \in \mathbb{R}^{m \times k}$ is an orthonormal basis with $\operatorname{rank}(\mathbf{H}^{\top}) = k, 1 \leq k \leq \min\{m, n\}$. The diagonal $\mathbf{S} := \operatorname{Diag}([\sigma_1, \ldots, \sigma_k]) \in \mathbb{R}^{k \times k}$ contains k singular values sorted in the descending order, where σ_l gives the percentage importance of \mathbf{u}_l . The corresponding singular vectors constitute $\mathbf{V} = [\mathbf{v}_1, \ldots, \mathbf{v}_k] \in \mathbb{R}^{n \times k}$.

Rather than using the full $\boldsymbol{U} \in \mathbb{R}^{m \times k}$, we only preserve the first *p*-columns of **U** ($p \le k$), denoted as **U**_{*p*}, to include the most important p components of the original space H. The subspace span(\boldsymbol{U}_p) composes a Grassmannian $\mathcal{U} := [\boldsymbol{U}_p]$ in $\mathcal{M}(p, m)$. In practice, it could potentially hurt the expressiveness of graph embedding for defining an identical relatively small subspace dimension p for all graphs, since real-world datasets could have a great number of graphs with a large variation on node sizes from a few to thousands. Instead, we let *p* for a Grassmann point $[\boldsymbol{U}_d] \in \mathcal{M}(p,m)$ be determined by $p = \sum_{i=1}^k \mathbb{1}\{\sigma_i > r\}$. The σ_i corresponds to unit singular values from (1), and r denotes the global threshold of the percentage importance. However, all these Grassmann points $\{\mathcal{U}\}$ can be naturally mapped to the embedded Grassmannian $\mathcal{M}(p_{\max}, m)$, with p_{\max} the highest p over all Grassmann points, so they are still at the same space with accessible geodesic distance.

The embedding operation is compatible to arbitrary hidden representations of \mathbb{H} . Intuitively, it regards a (sub)graph with a set of *n* attributed nodes as a *p*-dimensional subspace. While a variant of size *n* can be observed from different node sets, the rectified Grassmann points from the underlying embedding operation are projected to the same Grassmann manifold, where the geodesic distance of two points reflects the similarity of two sets, or graph instances in analogue. Such similarity provides a criterion for distance-based training tasks, such as clustering or classification.

The rectification step by the truncated SVD can be considered as a non-linear transformation of the feature space that extracts the most powerful subspace expression of the node space and view it as a Grassmann point. At this stage, node sets of varying sizes are embedded in a common Grassmann manifold, where each of them is represented by a subspace of orthonormal basis $U_p \in \mathbb{R}^{m \times p}$. While it is feasible to compute the geodesic distance of Grassmann points, projecting them back to the Euclidean space is preferred by conventional deep learning modules. We now introduce the projection operation of a Grassmann point to its associated Euclidean representation.

4.3. Projection embedding

The rectification step establishes a set of Grassmann points $\{U_1, \ldots, U_N\} \subset \mathcal{M}(p, m)$ as well as their matrix representation $\{U_1, \ldots, U_N\}$ from the orthonormal basis perspective. One can establish follow-up learning schemes on $\{U_1, \ldots, U_N\}$ with the Grassmann geometry (Absil et al., 2008). Alternatively, the Grassmann points can first be projected to the space of symmetric matrices Sym(*m*) to allow calculations based on Euclidean space. Define

$$\Pi : \mathcal{M}(p, m) \longrightarrow \operatorname{Sym}(m), \quad \Pi(\mathbf{U}) = \mathbf{U}\mathbf{U}^{\top}.$$
(2)

The projected representation Sym(m) allows general Euclidean measures to allow conventional deep learning methods such as fully-connected layers.

After the rectification and the projection step, the graph latent representation \boldsymbol{H} is transformed to a symmetric positive definite (SPD) matrix $\Pi(\boldsymbol{U}_p) = \boldsymbol{U}_p \boldsymbol{U}_p^{\top}$. This SPD matrix representation is an analog to a bilinear mapping, and it captures the second-order statistics that better reflects regional features of \boldsymbol{H} (Tuzel, Porikli, & Meer, 2006). Moreover, the rectified representation \boldsymbol{U} from (1) gains robustness as a result of approximating the covariance matrix $\boldsymbol{U}_p \boldsymbol{U}_p^{\top}$. This projected Euclidean representation is feasible for various tasks. For example, in graph property prediction, a vectorized $\Pi(\boldsymbol{U}_p)$ can be employed as the readout train. We summarize the main steps of the Grassmann embedding in Algorithm 1.

4.4. Stable SVD for backward propagation

While we package the embedding operations in an end-to-end learning framework, it is essential to develop a computational strategy for SVD that is reliable in back-propagation (BP) of deep neural networks. The rest of this section gives the derivation of BP for the employed truncated SVD, which is numerically stable especially in the case when the input matrix involves extremely small singular values. We denote two orthonormal matrices $\boldsymbol{U} \in \mathbb{R}^{m \times k}$, $\boldsymbol{V} \in \mathbb{R}^{n \times k}$, and $\boldsymbol{S} = \text{diag}(s_1, s_2, \ldots, s_k) \in \mathbb{R}^{k \times k}$ as the output of SVD on H^{\top} during forward-propagation. To update H^{\top} in BP, its gradient is calculated by

$$\nabla_{\boldsymbol{H}^{\top}} f = \left[\boldsymbol{U} \left(\boldsymbol{F} \circ \left[\boldsymbol{U}^{\top} \overline{\boldsymbol{U}} - \overline{\boldsymbol{U}}^{\top} \boldsymbol{U} \right] \right) \boldsymbol{S} + \left(\boldsymbol{I}_{m} - \boldsymbol{U} \boldsymbol{U}^{\top} \right) \overline{\boldsymbol{U}} \boldsymbol{S}^{-1} \right] \boldsymbol{V}^{\top} + \boldsymbol{U} \left[\boldsymbol{S} \left(\boldsymbol{F} \circ \left[\boldsymbol{V}^{\top} \overline{\boldsymbol{V}} - \overline{\boldsymbol{V}}^{\top} \boldsymbol{V} \right] \right) \boldsymbol{V}^{\top} + \boldsymbol{S}^{-1} \overline{\boldsymbol{V}}^{\top} \left(\boldsymbol{I}_{n} - \boldsymbol{V} \boldsymbol{V}^{\top} \right) \right] + \boldsymbol{U} \left(\boldsymbol{I}_{k} \circ \overline{\boldsymbol{S}} \right) \boldsymbol{V}^{\top},$$

where $\mathbf{F}_{ij} = \frac{1}{s_j^2 - s_i^2} \cdot \mathbb{1}\{i \neq j\}$ that satisfies the identity $\mathbf{F}^{\top} = -\mathbf{F}$. The calculations of \mathbf{F} and \mathbf{S}^{-1} are often numerically unstable due to the possible near-zero singular values. To circumvent this difficulty, we follow Huang and Van Gool (2017) and introduce the following trick on \mathbf{S} :

$$\boldsymbol{S}_{i,i}^{\text{new}} = \boldsymbol{S}_{i,i} \cdot \mathbb{1}\{\boldsymbol{S}_{i,i} > \epsilon\} + \epsilon \cdot \mathbb{1}\{\boldsymbol{S}_{i,i} \le \epsilon\},\$$

where ϵ is a small number and can usually be set to 10^{-12} . In practice, we replace *S* by the modified matrix *S*^{new} in BP to avoid 0 values in *S*.

5. Applications on graph

This section applies the proposed EGG to two distinct graph learning tasks: node clustering and graph classification. We start with formulating the two problems to be solved, following the designed model structure of the two addressed problems.

Algorithm 1: EGG for graph representation learning

Input : Hidden representation H, dimension p. **Output:** A graph representation.

- 1 (for graph-level tasks): Transpose **H** to \mathbf{H}^{\top} .
- **2** Find the *p*-dimensional low-rank representation of \mathcal{U} by (1).

5.1. Graph classification

A graph-level representation learning task, such as graph classification and regression, takes multiple graphs $\mathbb{G} = \{\mathcal{G}_1, \ldots, \mathcal{G}_N\}$ as the input to train a feasible model that makes correct assignment $\mathbf{y}_i = g(\mathbf{X}_i, \mathbf{A}_i)$. The node sizes of different graphs are merely identical, and it is the duty of graph pooling to learn from a hidden graph representation $\mathbf{H}_i = f_1(\mathbf{X}_i, \mathbf{A}_i)$ so that the graph is summarized to $\mathbf{h}_i = f_2(\mathbf{H}_i)$ with a determined length. The representation is later employed for label prediction, i.e., $\mathbf{y}_i = f_3(\mathbf{h}_i)$.

With GNNs, the first step of $f_1(\cdot)$ is usually executed by graph convolutional layers to extract structure and node features, which outputs a hidden representation $H \in \mathbb{R}^{n \times m}$ for an arbitrary graph of *n* nodes. A graph pooling strategy is then selected to design a proper $f_2(\cdot)$ that unifies the dimension of representations to all the graphs. The proposed EGG defines $f_2(\cdot)$ by embedding each *H* to a Grassmann point of $\mathcal{M}(p, m)$. Specifically, it calculates a *k*-dimensional subspace of H^{\top} by the truncated SVD. The *U* from (1) denotes a Grassmann point associated with a graph. To allow Euclidean computations, it can be projected to a symmetric matrix by (2) and be flattened and sent to a classifier, such as fully-connected layers. The embedding step is also illustrated in Fig. 1 and Algorithm 1.

The validity of EGG for graph pooling is guaranteed by the two essential requirements: *uni-size representation* and *permutation invariance*, which we shall check now.

Proposition 1. EGG always produces a graph embedding $g \in \mathbb{R}^{\frac{m(m+1)}{2}}$ for the node representation matrix $H \in \mathbb{R}^{n \times m}$, regardless of the graph size n.

Proof. Given the node representation matrix $\boldsymbol{H} \in \mathbb{R}^{n \times m}$ of a graph \mathcal{G} with *n* nodes and *m* features, the Grassmann graph embedding gives the output $\boldsymbol{U}\boldsymbol{U}^{\top}$, where $\boldsymbol{U} \in \mathbb{R}^{m \times k}$ with $k = \operatorname{rank}(\boldsymbol{H}^{\top})$. Then, the output of EGG for the graph classification is the flattened representation of the upper triangular matrix of $\boldsymbol{U}\boldsymbol{U}^{\top} \in \mathbb{R}^{m \times m}$. The length of vector representation is thus m(m + 1)/2, and it is independent of the graph size *n*. \Box

Proposition 2. EGG satisfies the requirement of permutation invariance so that it produces the same Grassmann graph embedding under row permutations of the input node representation matrix.

Proof. Suppose H_1 is the node representation matrix of a graph and let $H_2 = PH_1$, where **P** is a permutation matrix. Then,

$$\boldsymbol{H}_1^{\top} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\top}, \quad \boldsymbol{H}_2^{\top} = \boldsymbol{H}_1^{\top}\boldsymbol{P}^{\top} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\top}\boldsymbol{P}^{\top}.$$

The Grassmann point for both H_1 and H_2 can be accessed by the same matrix U. Hence, the proposed graph embedding method is permutation invariant. \Box

5.2. Node clustering

Node-level tasks make predictions on each node v of a single graph G. For node clustering, the target is to segment the full graph into a number of subsets, where nodes from the same

// Transpose // Manifold Rectification // Projection Embedding

subset usually have a closer connection to each other, and/or share similar properties.

A typical approach in literature to solve node clustering problems is by using *Variational Graph Auto-Encoder* (VGAE) (Kipf & Welling, 2016) to generate a latent representation for a graph and then applying *k*-means (Lloyd, 1982) for clustering the nodes, where the VGAE model pursues the optimal variational parameters **W** that minimize the variational lower bound

$$\mathcal{L} = \mathbb{E}_{q(\boldsymbol{Z}|\boldsymbol{X},\boldsymbol{A})} |\log p(\boldsymbol{A}|\boldsymbol{Z})] - \mathrm{KL}[q(\boldsymbol{Z}|\boldsymbol{X},\boldsymbol{A})|p(\boldsymbol{Z})|.$$

Here $q(\mathbf{Z}|\mathbf{X}, \mathbf{A})$ is the encoder such that

$$q(\boldsymbol{Z}|\boldsymbol{X},\boldsymbol{A}) = \prod_{i=1}^{N} q(\boldsymbol{z}_i|\boldsymbol{X},\boldsymbol{A}), \quad q(\boldsymbol{z}_i|\boldsymbol{X},\boldsymbol{A}) \sim \mathcal{N}(\boldsymbol{z}_i|\boldsymbol{\mu}_i, \operatorname{diag}(\boldsymbol{\sigma}_i^2)).$$

Both the mean μ and log standard deviation log σ are approximated by graph convolutional layers, such as GCN (Kipf & Welling, 2017). The decoder is defined as

$$p(\boldsymbol{A}|\boldsymbol{Z}) = \prod_{i=1}^{N} \prod_{j=1}^{N} p(A_{ij}|\boldsymbol{z}_i, \boldsymbol{z}_j), \quad p(A_{ij} = 1|\boldsymbol{z}_i, \boldsymbol{z}_j) = \sigma(\boldsymbol{z}_i^{\top} \boldsymbol{z}_j),$$

where **A** is with respect to the adjacency matrix, and $\sigma(\cdot)$ is the activation function. We take the inference set $\mathbf{H} := [\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_m] \in \mathbb{R}^{n \times m}$ as the hidden representation of the graph with *n* nodes of hidden size *m* and send it to EGG for Grassmann embedding. The consequent $\mathbf{U}\mathbf{U}^\top \in \mathbb{R}^{n \times n}$ is handled by clustering methods, such as *k*-means (Lloyd, 1982), to assign clusters with *k* set to be the number of classes of the dataset. Instead of relying on the pairwise connection or the feature-space information, EGG leverages second-order correlations of nodes for a proper segmentation.

6. Experiment

This section evaluates the proposed framework on graph-level classification tasks as well as node-level clustering tasks. The former is conducted on six benchmarks of variant graph sizes, volume, and density, and the latter employs five popular graph datasets of moderate volume. All benchmark datasets and base-line methods are publicly available in the PyTorch Geometric (PYG) (Fey & Lenssen, 2019) library. The implementation of EGG is published at https://github.com/conf20/Egg. The rest of this section lists the experimental setup and analyzes the performance comparison of the two experiments.

6.1. Ablation study on graph pooling

6.1.1. Setup

We benchmark the performance on five binary classification and one multi-class classification tasks. The **Molhiv** (Hu et al., 2020) is from open graph benchmark, and all other datasets are provided by **TUDataset benchmarks** (Morris et al., 2020). The summary statistics of the six benchmark datasets is provided in Table 1. Most benchmark datasets preserve their original feature attributes except for the feature-less dataset **Collab**, where we follow Xu et al. (2019) to generate new features with one-hot encoding of node degrees. Also, virtual nodes (Ishiguro, Maeda,

³ Project \mathcal{U} to an SPD matrix (2).

Table 1

Summary of the datasets for graph property prediction tasks

, , , , , , , , , , , , , , , , , , ,	011	I J I I I I I I I I I I I I I I I I I I				
Datasets	Proteins	D&D	NCI1	Mutagen	Collab	Molhiv
# Graphs	1,113	1,178	4,110	4,337	5,000	41,127
# Classes	2	2	2	2	3	2
Min # nodes	4	30	3	30	32	2
Max # nodes	620	5,748	111	417	492	222
Avg # nodes	39	284	30	30	74	26
Avg # edges	73	716	32	31	2,458	28
# Features	3	89	37	14	0	9

Table 2

Performance comparison for graph classification with GCN convolution.

	Proteins	D&D	NCI1	Mutagen	Collab	Molhiv
TopKPool	73.48±3.57	74.87±4.12	75.11±3.45	79.84±2.46	81.18±0.89	77.11±1.27
SAGPool	75.89±2.91	74.96±3.60	76.30±1.53	79.86±2.36	79.26±5.37	75.36±1.82
EDGEPool	75.60±2.40	67.60±0.51	77.17±1.49	70.34 ± 1.69	75.09±0.81	75.14±1.66
PANPool	72.41±3.58	$72.52{\scriptstyle\pm4.05}$	$62.82{\scriptstyle\pm3.04}$	$70.14{\scriptstyle \pm 1.85}$	$75.78{\scriptstyle\pm2.03}$	$74.18{\scriptstyle\pm1.52}$
SUM	74.91±4.08	78.91±3.37	76.96±1.70	80.69±3.26	80.76±1.56	74.88±2.64
AVG	73.13±3.18	76.89±2.23	73.70±2.55	80.37±2.44	81.24±1.34	77.69±1.17
MAX	73.57±3.94	75.80±4.11	75.96±1.82	78.83±1.70	82.28±2.10	76.95±0.94
Attention	73.93±5.37	$77.48{\scriptstyle\pm2.65}$	$74.04{\scriptstyle\pm1.27}$	80.25±2.22	$81.58{\scriptstyle\pm1.72}$	77.44±1.27
EGG	77.79±2.16	79.10±2.98	77.80±2.01	81.01±1.28	82.94±1.06	76.60±1.10

Table 3

Performance comparison for graph classification with GIN convolution.

	Proteins	D&D	NCI1	Mutagen	Collab	Molhiv
TopKPool	73.66±6.00	76.40±2.32	77.06±0.90	78.30±1.36	$81.40{\scriptstyle\pm0.94}$	78.14±0.62
SAGPool	75.95±4.52	68.94±7.62	76.97±2.94	78.86±1.58	81.76±1.57	75.26±2.29
EDGEPool	75.13±3.62	72.82 ± 1.40	77.79±2.80	81.01±0.82	79.20±1.66	75.30±2.01
PANPool	71.43±2.15	$72.75{\scriptstyle\pm2.32}$	$71.68{\pm}4.45$	$78.09{\scriptstyle\pm1.27}$	$80.22{\scriptstyle\pm2.02}$	77.18 ± 1.13
SUM	78.04±2.30	78.57±1.26	78.83±1.49	81.31±1.10	82.64±0.85	77.41±1.16
AVG	71.70±2.08	74.37 ± 1.32	76.55±1.72	80.97±1.18	83.30±0.77	78.21±0.90
MAX	76.70±1.57	77.31±2.06	79.27±1.38	80.28±0.83	80.94±0.72	78.16±1.33
Attention	$75.63{\scriptstyle\pm1.13}$	71.76±3.26	$78.22{\scriptstyle\pm1.32}$	78.54±5.37	$83.22{\pm}0.30$	74.44 ± 2.12
EGG	79.80 ±1.09	81.18±1.14	81.31±1.55	82.53 ±0.72	81.32±0.68	77.82±0.90

& Koyama, 2019) are included in **Molhiv** to enhance the learning ability, as is suggested by Hu et al. (2020).

To learn the hidden representation of graph topological embedding for pooling layers, we consider two variants of GCN (Kipf & Welling, 2017) and GIN (Xu et al., 2019). For the first five datasets from **TU Dataset**s, we construct three GCN layers for **NCI1**, and two GCN layers for the other four datasets to encode graph hidden representation for pooling. Meanwhile, the number of GIN convolutional layers is set to four with the JKNET (Xu et al., 2018) construction. Both models for the largest dataset **Molhiv** use four convolutional layers. The learned hidden representation is sent to one of the baselines, following a two-layer MLP with size 64 and 16, respectively.

A fair comparison of EGG is made against four hierarchical and four global pooling methods. The former includes TOP-KPOOL (Cătălina et al., 2018; Gao & Ji, 2019), SAGPOOL (Lee et al., 2019) EDGEPOOL (Diehl, 2019; Diehl, Brunner, Le, & Knoll, 2019), and PANPOOL (Ma, Xuan, Wang, Li, & Liò, 2020), and the latter chooses ATTENTION (Li et al., 2016), SUMMATION, AVERAGE and MAXIMIZATION methods. The official implementation are of all eight baselines are provided by PyTorch Geometric (Fey & Lenssen, 2019).

The models are trained on 80% randomly selected samples of the datasets, validated on 10% samples, and tested on the left 10% samples. The main hyper-parameters are fine-tuned with a grid search engine, where we are interested in Learning rate in $\{5e-3, 1e-3, 5e-4\}$, L_2 weight decay in $\{5e-3, 1e-3, 5e-4\}$, and hidden units in $\{32, 64\}$ for the convolution layers. For the fully-connected layer, we search in particular the dropout ratio in $\{0, 0.5\}$. For EGG, we include an extra hyper-parameter of

the threshold ratio r in truncated SVD, which is set to one of $\{0.5, 0.8\}$. The model stops training whenever the validation loss stops improving for 20 consecutive epochs, or the training epoch reaches 200.

6.1.2. Result

Table 2 and Table 3 compare the prediction performance of EGG with the eight baselines. Follow the convention, we report the percentage value of mean test accuracy for the classification tasks with **TUDataset**s and ROC-AUC score for **Molhiv**. The mean performance score are averaged over 10 repetitions with their standard deviations provided after the \pm signs. In general, our EGG achieves the top score on all the six tasks with a lower volatility. The advantage is more salient when the hidden representation of graphs are embedded by GIN convolutions with JKNET structure. EGG constructs a feasible global pooling method that interprets the second-order correlation of the compressed graph expressions, which are more informative than the firstorder relationships. As this covariance relationship is captured by a non-linear transformation, i.e., a truncated SVD, one layer of EGG is generally sufficient for graph distilling tasks, and it reliefs the burden of the fine-tuning work in training a deep graph learning model. It should be emphasized that the correlation analysis of EGG relies on the node attributes. When the input graph is feature-less, such as Collab, the performance can be less-promising.

6.2. Node clustering

6.2.1. Setup

The second experiment validates the design of EGG to node clustering tasks. Five popular benchmarks, including the three

Table 4

Summary of the datasets for node clustering tasks.

	Cora	Citeseer	Pubmed	Wiki-CS	Coauthor-CS
# Nodes	2,708	3,327	19,717	11,701	18,333
# Edges	5,429	4,732	44,338	216,123	100,227
# Features	1,433	3,703	500	300	6,805
# Clusters	7	6	3	10	15
π clusters	/	0	J	10	

citation networks (**Cora, Citeseer, Pubmed**) (Yang, Cohen, & Salakhudinov, 2016), **Wiki-CS** (Mernyei & Cangea, 2020) and **Coauthor-CS** (Shchur, Mumme, Bojchevski, & Günnemann, 2018), are employed to examine the effectiveness of an additional embedding step of EGG. The statistics of the five datasets are summarized in Table 4. In terms of the baseline method, we train a VGAE model (Kipf & Welling, 2016) to generate the latent representation $H \in \mathbb{R}^{n \times m}$ of the graph and then send the H matrix to *k*-means (Lloyd, 1982) for clustering. Based on this baseline architecture, we insert our EGG before the *k*-means as an additional step. Specifically, we first send the learned latent representation H to EGG, then perform *k*-means on the output of the EGG procedure.

The same structure of the VGAE model is adopted from the work by Kipf and Welling (2016), where the encoder is implemented by a two-layer GCN (Kipf & Welling, 2017) and the decoder is simply given by an inner product between the latent variables as illustrated in Section 5.2. We train the VGAE model for 200 epochs using ADAM (Kingma & Ba, 2015) with a learning rate of 0.01. The dimensions of the hidden layer and the latent space are set to 32 and 16 in all experiments, respectively. To obtain the latent representation of nodes, VGAE is trained on a link prediction task of identifying edges and non-edges. Following Kipf and Welling (2016), we divide the dataset into the training, validation, and test sets by a random selection of 85%, 5%, and 10% of the total node connections. The positive and negative edges share the same amount. In terms of evaluation, we employ the following four metrics to validate the clustering results: Accuracy (Acc.), Normalized Mutual Information (NMI), Average Rand Index (ARI), and Completeness Score (CS).

6.2.2. Result

Table 5 reports the experimental results of EGG-enhanced *k*-means with VGAE in node clustering tasks. In addition to the four evaluation metrics, we also attach two additional ratios in the name of *k*-means+EGG_x (*y*), where the first ratio $x \in (0, 1)$ is used to determine the number of most important components *p* in the latent space of $\mathbf{H} \in \mathbb{R}^{n \times m}$ we have kept, that is $p = \lceil xm \rceil$. The second value *y* indicates the percentage of the information from the latent representation \mathbf{H} being captured by these *p* components.

It shows clearly from Table 5 that the EGG0.8-enhanced *k*-means consistently outperforms the plain *k*-means on all the datasets in terms of all four metrics. We also observe that EGG0.5 provides a comparable performance against the baseline, giving that the first $p = \lceil 0.5m \rceil$ components contain 90% variations of the graphs' latent representation. These observations confirm that EGG brings solid performance gains to the node clustering tasks.

7. Further investigation

This section validates EGG from three perspectives. The sensitivity of the model performance to the newly introduced hyperparameter is tested in the first part. We then explore the expressiveness of the learned embedding with a t-SNE visualization of the learned hidden representation. In the last sector, we check the learning behavior of EGG through the loss curve of training tasks.



Fig. 2. Sensitivity analysis for the threshold information ratio r on PROTEINS.

7.1. Sensitivity to the embedded dimension

As discussed in Section 4, the number of the subspace dimension *p* is adaptively selected for each Grassmann point, according to the global threshold of the percentage importance r. The ris thus a new hyperparameter for EGG. This section designs a sensitivity analysis to demonstrate the negligible impact of the choice of *r* on the performance of the Grassmann embedding. The model learns a graph classification task on **Proteins** with both GCN and GIN networks, which architectures and training setups are detailed in Section 6.1. We report the mean test accuracy over 10 repetitive runs. Seven different values of the threshold ratio r is set from 0.3 to 0.9 with step size 0.1. The results are visualized in Fig. 2, which draws a nearly horizontal trend of the mean test accuracy movement of EGG with different choices of the threshold ratio r. This suggests that the considerably wide choice of the hyperparameter r does not drastically influence the performance of EGG pooling. In fact, we suggest a moderately high value of r, such as 0.5, to retain the essential information of a graph in Grassmann embedding and guarantee a relatively fast computational speed of the algorithm at the same time.

7.2. Embedding expressiveness

Next, we exploit the expressiveness of the flattened Euclidean graph embeddings with two-dimensional t-distributed Stochastic Neighbor Embedding (t-SNE). The results are from the 3-class graph classification task **Collab** that we conducted in the first experiment of Section 6. In Fig. 3, each point denotes a graph hidden representation by EGG, and the three colors indicate one of the three true labels. For a more clear presentation, we sample 4000 instances in random. In the case of the GIN convolution, outputs from all the four pooling layers are aggregated, due to the employed JKNET structure. Both Figs. 3(b) and 3(c) suggest a clear clustering pattern of the pooled graphs.

A similar visualization for the learned representation in node clustering tasks is displayed in Fig. 4, where the three citation networks are trained. Once again, points of different colors are basically located at distinct corners, which implies that the hidden representations from EGG(0.8)+k-means manage to collect a differentiable pattern for the underlying clustering task.

7.3. Learning behaviors

In the last investigation, we check the training and validation curves for loss and accuracy in Fig. 5. The results are retrieved from the graph classification learning task on **Proteins** with the same experimental settings in Section 6.1. Here only a single run results are retrieved due to the employment of the early stopping criteria, with which every independent run could stop at a different epoch. Except for the minor volatility after epoch 5

Table 5

Performance comparison for node clustering. All scores are averaged over 10 repetitions with the scale between 0 and 1. The value after \pm is standard deviation.

	Method	Acc.	NMI	ARI	CS
	k-means	0.6137 ± 0.0345	0.4459 ± 0.0190	0.3782 ± 0.0312	0.4351±0.0188
<u>n</u>	k-means+EGG0.2 (53.16%)	0.5207±0.0553	0.3789 ± 0.0374	0.2967±0.0475	0.3702 ± 0.0360
S	k-means+EGG0.5 (90.73%)	0.6195 ± 0.0340	0.4345 ± 0.0269	0.3804 ± 0.0434	0.4292 ± 0.0257
	k-means+EGG0.8 (98.76%)	0.6388±0.0386	0.4548±0.0158	0.3998±0.0293	0.4591±0.0191
	k-means	0.4347 ± 0.0341	0.1931±0.0343	0.1548±0.0262	0.1942±0.0358
see	k-means+EGG0.2 (45.28%)	0.4156 ± 0.0409	0.1794 ± 0.0259	0.1527 ± 0.0294	0.1783 ± 0.0254
ite	k-means+EGG0.5 (88.25%)	0.4683±0.0323	0.2098±0.0283	0.1865±0.0301	0.2096 ± 0.0282
0	k-means+EGG0.8 (97.98%)	0.4698±0.0284	$0.2071{\pm}0.0235$	$0.1859{\scriptstyle\pm0.0282}$	0.2107±0.0201
-	k-means	0.6469±0.0232	0.2425±0.0233	0.2380±0.0368	0.2393±0.0233
ne	k-means+EGG0.2 (64.98%)	0.6339 ± 0.0193	0.2311 ± 0.0162	0.2196 ± 0.0228	0.2280 ± 0.0163
Idu	k-means+EGG0.5 (97.18%)	0.6296 ± 0.0228	0.2394 ± 0.0252	0.2171 ± 0.0302	0.2377 ± 0.0263
Ы	k-means+EGG0.8 (99.48%)	0.6521±0.0135	0.2532 ± 0.0201	0.2436±0.0231	0.2521±0.0224
S	k-means	0.4080 ± 0.0371	0.3429±0.0157	0.2156 ± 0.0240	0.3418±0.0153
Ŭ.	k-means+EGG0.2 (52.48%)	0.3875±0.0392	0.3207 ± 0.0183	0.2065 ± 0.0276	0.3195 ± 0.0179
/ik	k-means+EGG0.5 (93.21%)	0.4103±0.0253	0.3319 ± 0.0132	0.2232 ± 0.0249	0.3307 ± 0.0128
\$	k-means+EGG0.8 (99.43%)	0.4685 ± 0.0210	0.3643±0.0193	0.2423±0.0223	0.3633±0.0181
S	k-means	0.6410±0.0171	0.6950±0.0311	0.5442±0.0427	0.6880±0.0300
)-JC	k-means+EGG0.2 (45.05%)	0.5483±0.0250	0.6095±0.0210	0.4472 ± 0.0269	0.6050±0.0235
the	k-means+EGG0.5 (92.67%)	0.5878 ± 0.0188	0.6750±0.0150	0.5033±0.0190	0.6735±0.0157
Dau	k-means+EGG0.8 (99.68%)	0.6700±0.0201	0.7248±0.0195	0.6150±0.0205	0.7230±0.0193



Fig. 3. The t-SNE visualizations of graph representations produced by EGG pooling on COLLAB with GCN before training; with GCN after training; and with GIN after training.



Fig. 4. The t-SNE visualizations of node representations with EGG(0.8)+k-means on Cora, Citeseer and Pubmed.

that is very likely brought about by stochastic gradients, all the four training lines validate an efficient convergence of EGG, where the loss curve stabilizes quickly after a few epochs. The training process is slightly longer for GIN convolution, which is partly due to the more sophisticated network architecture for the model to fit.

8. Discussion and conclusion

This paper develops a Grassmann geometry-based graph embedding strategy named EGG. For a given set of hidden feature subspace of graphs, the proposed method rectifies them to Grassmann points of a Grassmann manifold to make analysis on them. Through establishing the view of treating graph nodes as a subspace, many new perspectives on formulating informative graph representations become visible. For example, this work approximates the covariance relationship of node attributes with non-linearity transformation, which concurrently offsets the inefficiency of a stack of fully-connected layers and blurs the minor perturbations of the initial representation. Furthermore, the new framework allows a swift projection from the manifold space back to the Euclidean space, so that the new representation supports common loss designs by Euclidean metrics. We demonstrate the effectiveness of the embedding framework with extensive numerical experiments, for both graph-level and node-level representation learning tasks.



Fig. 5. Training and Validation learning curves on the Proteins dataset with EGG pooling. The graph convolutional layers are set to GCN and GIN, respectively.

The proposed EGG has multiple aspects of potential. As said, it defines a non-linear transformation routine to graph features, which can be exceptional when working on complex or massive attributes. Moreover, treating regional graphs as Grassmann points or other entities from non-Euclidean space brings extra flexibility to the model design. We believe this idea could motivate more links of graph topological learning and other geometric learning schemes.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors thank the associated editor and reviewers for their constructive comments and suggestions. Ming Li acknowledges support from the National Natural Science Foundation of China (No. 62172370, No. U21A20473).

References

- Absil, P.-A., Mahony, R., & Sepulchre, R. (2004). Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. Acta Applicandae Mathematica, 80(2), 199–220.
- Absil, P.-A., Mahony, R., & Sepulchre, R. (2008). Optimization algorithms on matrix manifolds. Princeton University Press.
- Bendokat, T., Zimmermann, R., & Absil, P.-A. (2020). A Grassmann manifold handbook: Basic geometry and computational aspects. arXiv:2011.13699.

- Bo, D., Wang, X., Shi, C., & Shen, H. (2021). Beyond low-frequency information in graph convolutional networks. In AAAI, vol. 35, no. 5 (pp. 3950–3957).
- Boumal, N., & Absil, P.-A. (2015). Low-rank matrix completion via preconditioned optimization on the Grassmann manifold. *Linear Algebra and its Applications*, 475, 200–239.
- Brockschmidt, M. (2020). Gnn-film: Graph neural networks with feature-wise linear modulation. In *ICML* (pp. 1144–1152).
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. (2017). Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4), 18–42.
- Bruna, J., Zaremba, W., Szlam, A., & Lecun, Y. (2014). Spectral networks and locally connected networks on graphs. In *ICLR*.
- Chikuse, Y. (2003). Statistics on special manifolds, vol. 174. Springer Science & Business Media.
- Cătălina, C., Veličković, P., Jovanović, N., Kipf, T., & Liò, P. (2018). Towards sparse hierarchical graph classifiers. In *NeurIPS workshop on relational representation learning*.
- Dai, W., Kerman, E., & Milenkovic, O. (2012). A geometric approach to low-rank matrix completion. *IEEE Transactions on Information Theory*, 58(1), 237–247.
- Diehl, F. (2019). Edge contraction pooling for graph neural networks. arXiv: 1905.10990.
- Diehl, F., Brunner, T., Le, M. T., & Knoll, A. (2019). Towards graph pooling by edge contraction. In *ICML 2019 workshop on learning and reasoning with graph-structured data*.
- Dong, X., Frossard, P., Vandergheynst, P., & Nefedov, N. (2014). Clustering on multi-layer graphs via subspace analysis on grassmann manifolds. *IEEE Transactions on Signal Processing*, 62(4), 905–918.
- Edelman, A., Arias, T. A., & Smith, S. T. (1998). The geometry of algorithms with orthogonality constraints. SIAM Journal on Matrix Analysis and Applications, 20(2), 303–353.
- Fey, M., & Lenssen, J. E. (2019). Fast graph representation learning with PyTorch Geometric. In ICLR workshop on representation learning on graphs and manifolds.
- Gallivan, K. A., Srivastava, A., Liu, X., & Van Dooren, P. (2003). Efficient algorithms for inferences on Grassmann manifolds. In *IEEE workshop on statistical signal* processing, 2003 (pp. 315–318). IEEE.
- Gao, H., & Ji, S. (2019). Graph U-nets. In ICML.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *NeurIPS*, vol. 70 (pp. 1263–1272).

Hamilton, W. L. (2020). Graph representation learning. Morgan & Claypool Publishers

- Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. In NIPS.
- Harandi, M. T., Salzmann, M., Jayasumana, S., Hartley, R., & Li, H. (2014). Expanding the family of Grassmannian kernels: An embedding perspective. In ECCV (pp. 408–423). Springer.
- He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., & Wang, M. (2020). Lightgcn: Simplifying and powering graph convolution network for recommendation. In Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval (pp. 639–648).
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., et al. (2020). Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*.
- Huang, Z., & Van Gool, L. (2017). A Riemannian network for SPD matrix learning. In AAAI, vol. 31.
- Huang, Z., Wang, R., Shan, S., & Chen, X. (2015). Projection metric learning on Grassmann manifold with application to video based face recognition. In *CVPR* (pp. 140–149).
- Huang, Z., Wu, J., & Van Gool, L. (2018). Building deep networks on Grassmann manifolds. In AAAI, vol. 32.
- Ishiguro, K., Maeda, S.-i., & Koyama, M. (2019). Graph warp module: An auxiliary module for boosting the power of graph neural networks. arXiv:1902.01020.
- Kim, D., & Oh, A. (2021). How to find your friendly neighborhood: Graph attention design with self-supervision. In *ICLR*.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *ICLR*.
- Kipf, T. N., & Welling, M. (2016). Variational graph auto-encoders. In NIPS workshop on Bayesian deep learning.
- Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Knyazev, B., Taylor, G. W., & Amer, M. (2019). Understanding attention and generalization in graph neural networks. In *NeurIPS*, vol. 32.
- Koch, O., & Lubich, C. (2007). Dynamical low-rank approximation. SIAM Journal on Matrix Analysis and Applications, 29(2), 434–454.
- Kolouri, S., Naderializadeh, N., Rohde, G. K., & Hoffmann, H. (2021). Wasserstein embedding for graph learning. In ICLR.
- Lee, J., Lee, I., & Kang, J. (2019). Self-attention graph pooling. In ICML.
- Li, Y., Tarlow, D., Brockschmidt, M., & Zemel, R. (2016). Gated graph sequence neural networks. In *ICLR*.
- Lin, T.-Y., RoyChowdhury, A., & Maji, S. (2015). Bilinear cnn models for fine-grained visual recognition. In *ICCV* (pp. 1449–1457).
- Lloyd, S. (1982). Least squares quantization in PCM. IEEE Transactions on Information Theory, 28(2), 129–137.
- Lui, Y. M. (2012). Advances in matrix manifolds for computer vision. Image and Vision Computing, 30(6–7), 380–388.
- Ma, Y., Wang, S., Aggarwal, C. C., & Tang, J. (2019). Graph convolutional networks with eigenpooling. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 723–731).
- Ma, Z., Xuan, J., Wang, Y. G., Li, M., & Liò, P. (2020). Path integral based convolution and pooling for graph neural networks. In *NeurIPS*, vol. 33 (pp. 16421–16433).
- Mémoli, F. (2011). Gromov-Wasserstein distances and the metric approach to object matching. Foundations of Computational Mathematics, 11, 417–487.
- Mernyei, P., & Cangea, C. (2020). Wiki-CS: A wikipedia-based benchmark for graph neural networks. In ICML workshop on graph representation learning and beyond workshop.

- Mesquita, D., Souza, A. H., & Kaski, S. (2020). Rethinking pooling in graph neural networks. In *NeurIPS*.
- Minh, H. Q., Murino, V., & Minh, H. Q. (2016). Algorithmic advances in Riemannian geometry and applications. Springer.
- Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., & Neumann, M. (2020). TUDataset: A collection of benchmark datasets for learning with graphs. In ICML workshop on graph representation learning and beyond.
- Ngo, T., & Saad, Y. (2012). Scaled gradients on grassmann manifolds for matrix completion. In NIPS (pp. 1412–1420).
- Shchur, O., Mumme, M., Bojchevski, A., & Günnemann, S. (2018). Pitfalls of graph neural network evaluation. In NeurIPS workshop on relational representation learning workshop.
- Slama, R., Wannous, H., Daoudi, M., & Srivastava, A. (2015). Accurate 3D action recognition using learning on the Grassmann manifold. *Pattern Recognition*, 48(2), 556–567.
- Tailor, S. A., Opolka, F., Lio, P., & Lane, N. D. (2021). Do we need anisotropic graph neural networks? In *ICLR*.
- Tuzel, O., Porikli, F., & Meer, P. (2006). Region covariance: A fast descriptor for detection and classification. In ECCV (pp. 589–600). Springer.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. In ICLR.
- Wang, B., Hu, Y., Gao, J., Sun, Y., Chen, H., Ali, M., et al. (2017). Locality preserving projections for grassmann manifold. In *IJCAI* (pp. 2893–2900).
- Wang, B., Hu, Y., Gao, J., Sun, Y., & Yin, B. (2014). Low rank representation on grassmann manifolds. In Asian conference on computer vision.
- Wang, Z., & Ji, S. (2020). Second-order pooling for graph neural networks. IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Wang, X., Ji, H., Shi, C., Wang, B., Ye, Y., Cui, P., et al. (2019). Heterogeneous graph attention network. In WWW (pp. 2022–2032).
- Wang, Y., Long, M., Wang, J., & Yu, P. S. (2017). Spatiotemporal pyramid network for video action recognition. In CVPR (pp. 1529–1538).
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks* and Learning Systems, 32(1), 4–24.
- Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How powerful are graph neural networks? In ICLR.
- Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.-i., & Jegelka, S. (2018). Representation learning on graphs with jumping knowledge networks. In *ICML*.
- Yang, Z., Cohen, W., & Salakhudinov, R. (2016). Revisiting semi-supervised learning with graph embeddings. In ICML.
- Ye, K., & Lim, L.-H. (2016). Schubert varieties and distances between subspaces of different dimensions. SIAM Journal on Matrix Analysis and Applications, 37(3), 1176–1197.
- Zhang, Z., Cui, P., & Zhu, W. (2020). Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*.
- Zheng, X., Zhou, B., Gao, J., Wang, Y. G., Liò, P., Li, M., et al. (2021). How framelets enhance graph neural networks. In *ICML*.
- Zheng, X., Zhou, B., Li, M., Wang, Y. G., & Gao, J. (2020). MathNet: Haarlike wavelet multiresolution-analysis for graph representation and learning. arXiv:2007.11202.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., et al. (2020). Graph neural networks: a review of methods and applications. *AI Open*, *1*, 57–81.
- Zhou, B., Gao, J., Tran, M.-N., & Gerlach, R. (2021). Manifold optimization-assisted gaussian variational approximation. *Journal of Computational and Graphical Statistics*, 30(4), 946–957.
- Zhou, B., Zheng, X., Wang, Y. G., Li, M., & Gao, J. (2021). Grassmann graph embedding. In ICLR 2021 workshop on geometrical and topological representation learning.