

Synergistic Graph Fusion via Encoder Embedding

Cencheng Shen^{a,*}, Carey Priebe^b, Jonathan Larson^c, Ha Trinh^c

^aUniversity of Delaware, 19716, Newark, United States

^bJohns Hopkins University, 21218, Baltimore, United States

^cMicrosoft Research, 98052, Redmond, United States

Abstract

In this paper, we introduce a method called graph fusion embedding, designed for multi-graph embedding with shared vertex sets. Under the framework of supervised learning, our method exhibits a remarkable and highly desirable synergistic effect: for sufficiently large vertex size, the accuracy of vertex classification consistently benefits from the incorporation of additional graphs. We establish the mathematical foundation for the method, including the asymptotic convergence of the embedding, a sufficient condition for asymptotic optimal classification, and the proof of the synergistic effect for vertex classification. Our comprehensive simulations and real data experiments provide compelling evidence supporting the effectiveness of our proposed method, showcasing the pronounced synergistic effect for multiple graphs from disparate sources.

Keywords: Graph Fusion, Encoder Embedding, Supervised Learning

1. Introduction

Graphs have become increasingly prevalent in various real-world scenarios, including social networks, communication networks, webpage hyperlinks, and biological systems [9, 18, 3, 5, 39, 38]. In graph data, we have a set of vertices denoted as $\{v_i, i = 1, \dots, n\}$, and edges represented as $\{e_j, j = 1, \dots, s\}$ connecting these vertices. This information can be captured by an $n \times n$ adjacency matrix \mathbf{A} , where $\mathbf{A}(i, j) = 1$ indicates the existence of an edge between vertices i and j , and $\mathbf{A}(i, j) = 0$ indicates the absence of an edge.

Traditionally, the adjacency matrix is binary and primarily used for representing unweighted graphs. Its capacity extends to handling weighted graphs, where the elements of \mathbf{A} are assigned the corresponding edge weights. Moreover, graph representation can be used for any data via proper transformations. For instance, when vertex attributes are available, they can be transformed into a pairwise distance or kernel matrix, thus creating a general graph structure. This adaptability enables the incorporation of additional information from vertex attributes, making graph representation a powerful tool for data analysis that goes beyond traditional binary or weighted graphs.

*Corresponding author

Email addresses: shenc@udel.edu (Cencheng Shen), cep@jhu.edu (Carey Priebe), jolarso@microsoft.com (Jonathan Larson), trinhha@microsoft.com (Ha Trinh)

Graph embedding is a fundamental and versatile approach for analyzing and exploring graph data, encompassing a wide range of techniques such as spectral embedding [25, 36], graph convolutional neural networks [13, 40], node2vec [11, 14], among others. By projecting the vertices into a low-dimensional space while preserving the structural information of the graph, graph embedding yields vertex representations in Euclidean space that facilitate various downstream inference tasks, such as community detection [12, 43], vertex classification [21, 13], outlier detection [24, 1], etc.

This manuscript addresses a crucial and increasingly prevalent scenario in data analysis where we encounter multiple graphs $\{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_M\}$ sharing a common vertex set [19, 15, 31, 2]. This situation is becoming more frequent as data collection accelerates, and it arises in diverse fields such as social network analysis, neuroscience, and linguistics. For instance, one may want to analyze different social network connections for the same group of individuals, different measurements on the same brain regions, hyperlinks in articles written in different languages, or citation networks enriched with additional attributes. In such settings, it is essential to develop a method that can effectively integrate data from multiple sources. The ideal method should not only exhibit improved performance as more signals from different graphs are incorporated but also be robust against noisy or irrelevant data sources. By leveraging information from multiple graphs, we aim to achieve a richer and more comprehensive representation of the underlying structure and relationships within the data. This can lead to better insights and enhanced accuracy in various downstream tasks, such as vertex classification and link prediction.

Existing works on multiple-graph settings are relatively limited. Most of these investigations, particularly those with theoretical foundations, are based on spectral embedding. For instance, the Omnibus method [22, 16] constructs an $Mn \times Mn$ matrix by placing the input adjacency on the diagonals and computing the average of two adjacency matrices on each off-diagonal $n \times n$ sub-matrix. This matrix is then projected into an $n \times d$ matrix using spectral embedding. The more recent multiple adjacency spectral embedding (MASE) [2] method conducts individual spectral embedding for each adjacency matrix, concatenating the resulting embeddings followed by single value decomposition to obtain an $n \times d$ embedding. Another method, unfolded spectral embedding (USE) [27], concatenates the adjacency matrices into an $n \times Mn$ matrix, then applies spectral embedding to obtain an $n \times Md$ embedding. The main difference among these methods lies in how they concatenate the multiple graphs in a proper manner. As they all utilize spectral embedding, specifically singular value decomposition, they require the selection of the dimension d . While they have demonstrated certain theoretical properties, the theoretical investigations do not focus on the subsequent downstream task and how well it benefits from collecting more graphs.

To that end, in this paper we present a supervised multi-graph embedding approach called "graph fusion embedding," which builds upon the recently proposed graph encoder embedding [34, 35, 28]. The proposed method is designed to accommodate multiple-graph of distinct structures, including binary, weighted, or

general graphs via distance or kernel transformation of vertex attributes. A notable feature of the proposed method is the "synergistic effect" in supervised learning. This effect guarantees that, for sufficiently large vertex sizes, the accuracy of vertex classification consistently benefits from the inclusion of additional graphs and never degrades. We prove this synergistic effect under both the stochastic block model for binary graphs and a more general graph model. Additionally, we explore the asymptotic behavior of graph fusion embedding and establish a sufficient condition for it to achieve asymptotically perfect vertex classification.

To validate the superiority of our proposed method, we conduct extensive simulations and evaluations on real graph data. By applying the graph fusion embedding to various types of graphs, we showcase its effectiveness and versatility in different settings. Our method consistently outperforms existing techniques and demonstrates the synergistic effect, making it a robust and powerful tool for multi-graph analysis. The MATLAB code for the method and simulations are made available on Github¹. The appendix includes additional simulations validating the proposed methods, all benchmark results using competitor methods, as well as all theorem proofs.

2. Method

2.1. Main Algorithm

The graph fusion embedding takes as input M graphs with a common set of n vertices and a label vector representing K communities. For vertices with unknown labels, the corresponding entry in the label vector is set to 0, and such labels are not used in the embedding process.

- **Input:** The graphs $\{\mathbf{A}_m \in \mathbb{R}^{n \times n}, m = 1, \dots, M\}$ and a label vector $\mathbf{Y} \in \{0, 1, \dots, K\}^n$.
- **Step 1:** For each $k = 1, \dots, K$, compute the number of observations per-class

$$n_k = \sum_{i=1}^n I(\mathbf{Y}_i = k).$$

- **Step 2:** Compute the one-hot encoding matrix $\mathbf{W} \in \mathbb{R}^{n \times K}$ based on \mathbf{Y} , then normalize by the number of observations per-class. Specifically, for each vertex $i = 1, \dots, n$, we set

$$\mathbf{W}(i, k) = 1/n_k$$

if and only if $\mathbf{Y}_i = k$, and 0 otherwise. Note that vertices with unknown labels are unused and effectively zero, i.e., $\mathbf{W}(i, :)$ is a zero vector if $\mathbf{Y}_i = 0$.

- **Step 3:** For each graph $m = 1, \dots, M$, calculate the graph encoder embedding:

$$\mathbf{Z}_m = \mathbf{A}_m \times \mathbf{W}.$$

¹<https://github.com/cshen6/GraphEmd>

- **Step 4:** Let $\mathbf{Z}_m(i, \cdot)$ denote each row of \mathbf{Z}_m , and $\|\cdot\|$ be the Euclidean norm. For each i and each m where $\|\mathbf{Z}_m(i, \cdot)\| > 0$, normalize the non-zero row by the norm

$$\tilde{\mathbf{Z}}_m(i, \cdot) = \frac{\mathbf{Z}_m(i, \cdot)}{\|\mathbf{Z}_m(i, \cdot)\|},$$

then construct the row-concatenated embedding $\tilde{\mathbf{Z}}$ by setting

$$\tilde{\mathbf{Z}}(i, \cdot) = [\tilde{\mathbf{Z}}_1(i, \cdot), \tilde{\mathbf{Z}}_2(i, \cdot), \dots, \tilde{\mathbf{Z}}_M(i, \cdot)] \in \mathbf{R}^{MK}$$

for each vertex i .

- **Output:** The graph fusion embedding $\tilde{\mathbf{Z}} \in \mathbf{R}^{n \times MK}$.

2.2. Subsequent Classification

In the supervised setting, the label vector is always available. During steps 1 to 3 of the graph fusion embedding process, the labels of training vertices are used to construct the W matrix, while the labels of testing vertices are set to zero and remain unused throughout the embedding process. It is important to note that both the training and testing vertices are present in the graph, and as a result, the resulting embedding $\tilde{\mathbf{Z}}$ provides representations for all vertices in the graph. Let \mathbf{Y}_i be a random variable taking values in $1, \dots, K$ that represents the true label for any testing vertex i , and let $g(\cdot)$ denotes a classifier function. The true classification error can be denoted as

$$L = \text{Prob}(g(\tilde{\mathbf{Z}}(i, :)) \neq \mathbf{Y}_i).$$

Throughout the main paper, we employed the 5-nearest-neighbor classifier as our classifier of choice. The synergistic effect is also observed when using other standard classifiers, such as discriminant analysis or a simple two-layer neural network. These alternatives classifiers can be readily applied, yielding similar numerical outcomes, and are presented in the appendix. Furthermore, the main algorithm can be adapted to an unsupervised version in an iterative manner, similar to the single graph version in [33], which is also briefly discussed in the appendix.

2.3. Computational Complexity

The computation complexity of the graph fusion embedding is $O(nMk + \sum_{m=1}^M s_m)$, where s_m is the number of entries in each graph. Step 1 and 2 are shared across all graphs, and step 3 is carried out separately for each graph, allowing for parallelization and further reducing the running time to $O(nk + \max_{m=1, \dots, T} \{s_m\})$. Note that if $M = 1$, the fusion embedding is equivalent to the normalized graph encoder embedding for a single graph. Additionally, for sparse graphs, step 3 can be implemented in linear time using two simple edgelist operations.

Figure 1 provides a comparison of the running time between the graph fusion embedding and the unfolded spectral embedding. The algorithm demonstrates linear scalability with respect to the number of graphs, with the three-graph fusion embedding taking approximately three times longer than the single-graph embedding. Additionally, the graph fusion embedding is significantly more efficient than the spectral embedding, requiring only 0.2 seconds at $n = 30000$, whereas the spectral embedding takes over 4 minutes. In general, the encoder embedding outperforms any existing graph embedding approach and can process millions of edges in just a few seconds.

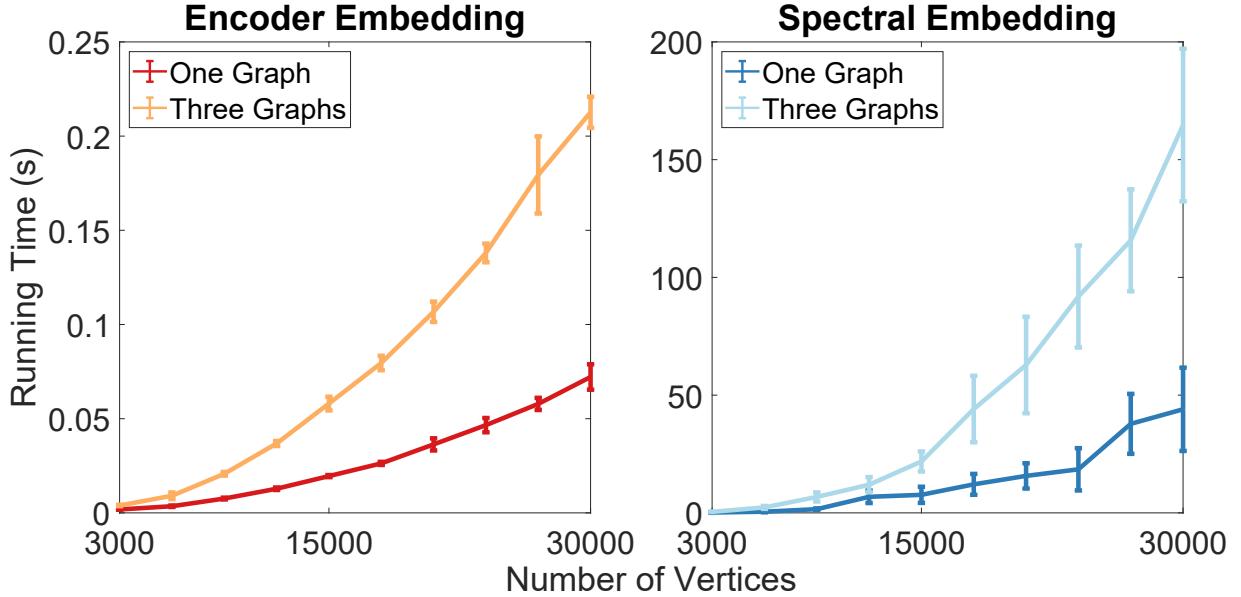


Figure 1: This figure illustrates the runtime performance of the graph fusion embedding (left panel) and the unfolded spectral embedding (right panel). The graphs are generated by simulation model 2 in Section 5, and we carry out 20 Monte-Carlo replicates at each sample size n . Each line represents the average running time, with error bars indicating the standard deviation.

3. Theory on Binary Graph

3.1. The Stochastic Block Model

The binary graph serves as the foundation of most graph methods and is commonly used in various applications. To demonstrate the functionality of graph fusion embedding on binary graphs, we focus on the stochastic block model (SBM) for both theoretical analysis and simulations.

Under SBM, each vertex i in a graph \mathbf{A}_m is assigned a label $\mathbf{Y}_i \in \{1, \dots, K\}$. The probability of an edge between a vertex from class k and a vertex from class l is determined by a block probability matrix

$\mathbf{B}_m = [\mathbf{B}_m(k, l)] \in [0, 1]^{K \times K}$, and for any $i \neq j$ it holds that

$$\mathbf{A}_m(i, j) \sim \text{Bernoulli}(\mathbf{B}_m(\mathbf{Y}_i, \mathbf{Y}_j)).$$

It is worth mentioning that this formulation generates a directed graph, but it can be easily adapted for an undirected graph by setting $\mathbf{A}_m(j, i) = \mathbf{A}_m(i, j)$ for all $i < j$. The choice of directed or undirected graph does not affect the method or theory, and in this case, we opt for the directed case. Additionally, the graph may or may not have self-loops, i.e., $\mathbf{A}_m(i, i)$ could be 0 or 1, and the method and the theorems still hold.

The degree-corrected stochastic block model (DC-SBM) is an extension of the SBM that accounts for the sparsity observed in real graphs [43]. In addition to the parameters defined in the SBM, each vertex i is associated with a degree $\theta_i \stackrel{i.i.d.}{\sim} F_\theta$. When conditioned on the degrees, each edge from vertex i to another vertex j is independently generated as follows:

$$\mathbf{A}_m(i, j) \sim \text{Bernoulli}(\theta_i \theta_j \mathbf{B}(\mathbf{Y}_i, \mathbf{Y}_j)).$$

As standard SBM is a special case where all degrees are 1, it suffices to consider the DC-SBM graphs for theory purpose.

Finally, the class labels \mathbf{Y}_i are assumed to follow a categorical distribution with prior probability $\{\pi_k \in (0, 1) \text{ with } \sum_{k=1}^K \pi_k = 1\}$. Under this assumption, as the number of vertices n increases to infinity, the number of vertices within each class n_k also converges to infinity. If not, π_k asymptotically approaches 0, rendering class k trivial in the model. This assumption implies that $n_k = O(n)$ for all k .

Note that we do not assume any specific dependence structure between graphs. For instance, $\mathbf{A}_1(i, j)$ and $\mathbf{A}_2(i, j)$ could be independently distributed, identical, similar up to some noise or transformation, or correlated in any other way, and all theorems would still hold true.

3.2. Convergence of the Fusion Embedding

Given M graphs from DC-SBM with a common vertex set, we define $\tilde{\mathbf{B}}_m$ as the row-normalized block probability matrix of \mathbf{B}_m , such that for each row k ,

$$\tilde{\mathbf{B}}_m(k, :) = \frac{\mathbf{B}_m(k, :)}{\|\mathbf{B}_m(k, :)\|}.$$

Then we form the row concatenation of all $\tilde{\mathbf{B}}_m$ as:

$$\tilde{\mathbf{B}} = [\tilde{\mathbf{B}}_1, \tilde{\mathbf{B}}_2, \dots, \tilde{\mathbf{B}}_M] \in \mathbb{R}^{K \times MK}.$$

The first theorem demonstrates that given DC-SBM graphs, the graph fusion embedding converges to the normalized and concatenated block probability matrix.

Theorem 1. Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. Let n be the number of vertices with known labels. Then, for any vertex i belonging to class \mathbf{Y}_i , its graph fusion embedding satisfies:

$$\|\tilde{\mathbf{Z}}(i, \cdot) - \tilde{\mathbf{B}}(\mathbf{Y}_i, :)\| = O\left(\frac{1}{\sqrt{n}}\right) \xrightarrow{n \rightarrow \infty} 0.$$

3.3. Asymptotic Perfect Separation

Under DC-SBM, the normalized and concatenated block probability matrix $\tilde{\mathbf{B}}$ fully determines the embedding. As a result, there exists a simple condition to ensure the embedding is asymptotically perfect for vertex classification. In the context of vertex classification, both the training and testing vertices follow the same asymptotic convergence in Theorem 1, with the only difference being that the training labels are known while the testing labels are not observed. We denote the finite-sample nearest-neighbor classification error using M graphs and n training vertices as $L_n(\mathbf{A}_1, \dots, \mathbf{A}_M)$.

Theorem 2. Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. As the number of training vertices increases to infinity, the graph fusion embedding achieves asymptotically perfect vertex classification, i.e.,

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) = 0,$$

if and only if there are no repeating rows in $\tilde{\mathbf{B}}$.

In essence, the graph fusion embedding can achieve perfect classification when there is no overlap between the rows of $\tilde{\mathbf{B}}$, meaning that each row represents a unique class. Conversely, if any two rows of $\tilde{\mathbf{B}}$ coincide, then the embedding vectors for vertices in those two classes will also coincide, making them indistinguishable.

While the condition may appear restrictive, it is equivalent to stating that every class of vertices must be perfectly separable from any other class, a scenario that rarely occurs in practice. Indeed, perfect classification serves as a preliminary step for the subsequent theorem on the synergistic effect, which always holds for our fusion embedding and does not require additional conditions.

3.4. The Synergistic Effect

Theorem 3. Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. For any subset of graphs $M_1 \leq M$, the classification error using the fusion embedding satisfies:

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) - L_n(\mathbf{A}_1, \dots, \mathbf{A}_{M_1}) \leq 0$$

for sufficiently large sample size n .

Therefore, as more graphs are collected and included, vertex classification via the graph fusion embedding always achieves equal or better performance. If an included graph contains pure noise, the classification error does not deteriorate; however, any graph that contains additional signal to better separate classes can lead to a reduction in classification error. It is important to note that the theorems are based on asymptotic analysis, and the finite-sample error may still exhibit minimal deterioration, especially for small vertex size.

4. Extension to General Graphs

4.1. A General Graph Model

To understand how the graph fusion embedding can be applied to weighted graphs or distance transformation of vertex attributes, we consider the following general graph model. For any fixed vertex i , we assume that

$$\mathbf{A}_m(i, 1), \mathbf{A}_m(i, 2), \dots \stackrel{i.i.d.}{\sim} F_m(i),$$

and the densities $\{F_m(i), m = 1, \dots, M \text{ and } i = 1, \dots, n\}$ have finite moments. Namely, each row of the adjacency matrix is independently and identically distributed when conditioning on vertex i . Note that when we condition on the vertex i , it means that all relevant information associated with that vertex are known, including its degree, class label, or Euclidean position, depending on the underlying distribution of the entries in the adjacency matrix.

This model includes a wide range of graph representations, including the aforementioned degree-corrected stochastic block model (DC-SBM), the random dot product graph model [42], as well as a pairwise distance or kernel matrix from Euclidean space. To illustrate, consider the scenario where we have vertex attributes X_i for each vertex $i = 1, \dots, n$, and we construct a Euclidean distance matrix \mathbf{A} where $\mathbf{A}(i, j) = \|X_i - X_j\|$. Under this transformation, when we condition on the attributes X_i , the entries $\mathbf{A}(i, j)$ and $\mathbf{A}(i, s)$ are independently and identically distributed for any $j \neq s$.

We reiterate that the general model does not assume any dependence structure between graphs. For instance, $\mathbf{A}_1(i, j)$ and $\mathbf{A}_2(i, j)$ could be correlated in any manner. The model assumption is solely on the edge generation within each individual graph.

4.2. Convergence under General Graph

Given the above model, we define the following conditional expectation for the m th graph:

$$E(\mathbf{A}_m(i, j)|i) = [E(\mathbf{A}_m(i, j)|i, \mathbf{Y}_j = 1), \dots, E(\mathbf{A}_m(i, j)|i, \mathbf{Y}_j = K)],$$

which is a K -dimensional vector. We then define the normalized vector as follows:

$$\mu_m(i) = \frac{E(\mathbf{A}_m(i, j)|i)}{\|E(\mathbf{A}_m(i, j)|i)\|} \in \mathbb{R}^K.$$

Finally, we form the concatenated mean matrix as follows:

$$\mu(i) = [\mu_1(i), \mu_2(i), \dots, \mu_M(i)] \in \mathbb{R}^{MK},$$

which represents the concatenation of all the normalized conditional expectation vectors for vertex i across all M graphs.

Theorem 4. *Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the general graph model in Section 4.1. Let n be the number of vertices with known labels. Then, for any vertex i belonging to class \mathbf{Y}_i , its graph fusion embedding satisfies:*

$$\|\tilde{\mathbf{Z}}(i, :) - \mu(i)\| = O\left(\frac{1}{\sqrt{n}}\right) \xrightarrow{n \rightarrow \infty} 0.$$

Thus, in the context of a general graph model, the fusion embedding provides n point masses represented by the concatenated conditional means. The location of each point mass is determined by the specific vertex i under consideration, reflecting the characteristics of the underlying graph model. For example, under the DC-SBM graph, $\mu(i) = \tilde{\mathbf{B}}(\mathbf{Y}_i, :)$, implying that all vertices belonging to the same class share a common point. This effectively reduces n embedding position to K points, where each point corresponds to a distinct class.

4.3. Vertex Classification under General Graph

Let $\{h_k(x), k = 1, \dots, K\}$ denotes the density family where each $h_k(x)$ represents all possible $\mu(i)$ from class k . Generally, $h_k(x)$ is continuous, but in specific cases, such as discrete graph models, the density can be discrete, e.g., under DC-SBM, $h_k(\tilde{\mathbf{B}}(k, :)) = 1$ and 0 otherwise. Using the density characterization, the next theorem expands the results of Theorem 2 and Theorem 3 to the general graph model. It establishes a general condition for achieving asymptotically optimal classification through the fusion embedding and demonstrates the presence of a synergistic effect in general scenarios.

Theorem 5. *Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the general graph model in Section 4.1. As the number of training vertices increases to infinity, the graph fusion embedding achieves asymptotically perfect vertex classification, i.e.,*

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) = 0$$

if and only if

$$\sum_{1 \leq k < l \leq K} h_k(x)h_l(x) = 0$$

for any $x \in \mathbb{R}^{MK}$.

Additionally, for any subset of graphs $M_1 \leq M$, the classification error using the fusion embedding satisfies:

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_{M_1}, \dots, \mathbf{A}_M) - L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) \leq 0$$

for sufficiently large sample size n .

Once more, the condition for asymptotically perfect classification implies that each class must be perfectly separable, as indicated by the class-wise density of the expectations. Although such perfect separability is not commonly observed, it serves as an important special case in understanding the synergistic effect.

Note that based on the convergence rate for the embedding ($O(\frac{1}{\sqrt{n}})$ as shown in Theorem 4), it is possible to estimate the required sample size to observe the classification improvement. However, the precise value of n is contingent not only on the convergence rate but also on factors like the magnitude of separation between classes and the standard deviation of the underlying graph distribution.

As an illustration, let us make some simplifying assumptions. Suppose each class-conditional mean (i.e., $\mu(i)$) is at least c_1 apart from others, and the standard deviation for the graph variable (i.e., $F_m(i)$) is c_2 . In this case, we can employ a normal approximation with three standard deviations, and deduce that the required sample size for classification error convergence must satisfy $6\frac{c_2}{\sqrt{n}} \geq c_1$.

For example, considering $c_1 = 0.1$ and $c_2 = 0.5$, solving the inequality yields $n \geq 900$, which is similar to the setting for the rightmost figure in Figure 3 with one graph. As another example, if we assume $c_1 = 0.02$ (indicating very small differences in class-conditional expectations) and $c_2 = 0.4$, then n would need to be approximately 15000 to achieve perfect classification.

5. Simulations

In this section, we present three simulations to illustrate the synergistic effect of the graph fusion embedding. We consistently use the 5-nearest-neighbor classifier in all simulations and real data experiments. While the main paper primarily focuses on the graph fusion embedding and the synergistic effect in classification, additional simulation results are provided in the appendix. These include experiments with heterogeneous graphs, alternative classifiers, the unsupervised case, and comparisons with benchmark methods (Omnibus, MASE, and USE).

Simulation 1

We consider a scenario with $K = 4$ classes, where the prior probabilities for each vertex to belong to class $[1, 2, 3, 4]$ are set to $[0, 0.3, 0.2, 0.2, 0.3]$. Three SBM graphs are generated with block probability matrices $B_m \in \mathbb{R}^{K \times K}$. For each m , the probabilities are $B_m(k, l) = 0.1$ for all k, l , except $B_m(m, m) = 0.2$ for $m = 1, 2, 3$. This means that the m th class has a stronger connection in graph m . A visualization of

this setup is provided in Figure 2. While no single graph is asymptotically perfect for classification, the concatenated block probability matrix

$$\mathbf{B} = \begin{bmatrix} 0.2 & \cdots & 0.1 & 0.1 & \cdots & 0.1 & 0.1 & 0.1 & \cdots \\ 0.1 & \cdots & 0.1 & 0.2 & \cdots & 0.1 & 0.1 & 0.1 & \cdots \\ 0.1 & \cdots & 0.1 & 0.1 & \cdots & 0.1 & 0.1 & 0.2 & \cdots \\ 0.1 & \cdots & 0.1 & 0.1 & \cdots & 0.1 & 0.1 & 0.1 & \cdots \end{bmatrix}$$

satisfies the condition in Theorem 2. Therefore, each graph provides additional information to aid in separating the communities, and using all three graphs leads to asymptotic perfect classification.

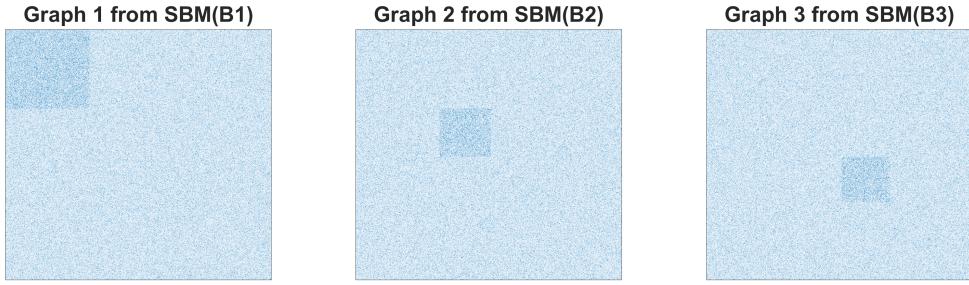


Figure 2: Visualize three SBM graphs at $n = 1000$.

Simulation 2

Simulations 2 uses the same settings as simulation 1, except that the three graphs are generated using DC-SBM. The parameters are the same as above, and the degrees are drawn independently from a uniform distribution in the interval $[0.1, 0.5]$.

Simulation 3

This simulation involves starting with one signal graph and progressively adding independent graphs. The number of classes remains at four with the same prior probability. We first generate \mathbf{A}_1 using SBM with the block probability matrix \mathbf{B}_1 , which has 0.1 in off-diagonal entries and 0.2 in diagonal entries. Next, we generate five independent SBM graphs \mathbf{A}_m for $m = 2, \dots, 6$ with the block probability matrix \mathbf{B}_m being 0.1 in all entries. The concatenated block probability matrix \mathbf{B} always satisfies Theorem 2, even in the presence of independent graphs, i.e.,

$$\mathbf{B} = \begin{bmatrix} 0.2 & 0.1 & 0.1 & 0.1 & \cdots \\ 0.1 & 0.2 & 0.1 & 0.1 & \cdots \\ 0.1 & 0.1 & 0.2 & 0.1 & \cdots \\ 0.1 & 0.1 & 0.1 & 0.2 & \cdots \end{bmatrix}$$

Theorems 2 and 3 predict that the classification error should stay almost the same with the addition of noisy graphs.

Results

In both simulation 1 and simulation 2, the classification error improves with the inclusion of more graphs, as illustrated in the left and center panels of Figure 3. Using two graphs yields a lower error than using only one graph, and using all three graphs yields the lowest error. This is because each new graph contains new signal about the labels, so including a new graph shall always yields better performance.

In simulation 3, only the first graph is related to the label, while all additional graphs are essentially noise. The addition of noise graphs has a minimal effect on the classification error of fusion embedding as the number of vertices increases, as shown in the right panels of Figure 3. Compared to the benchmark methods' performance in the appendix, the graph fusion embedding stands out as the only method that demonstrates the synergistic effect, remains robust against noise graphs, and consistently performs well across all simulations.

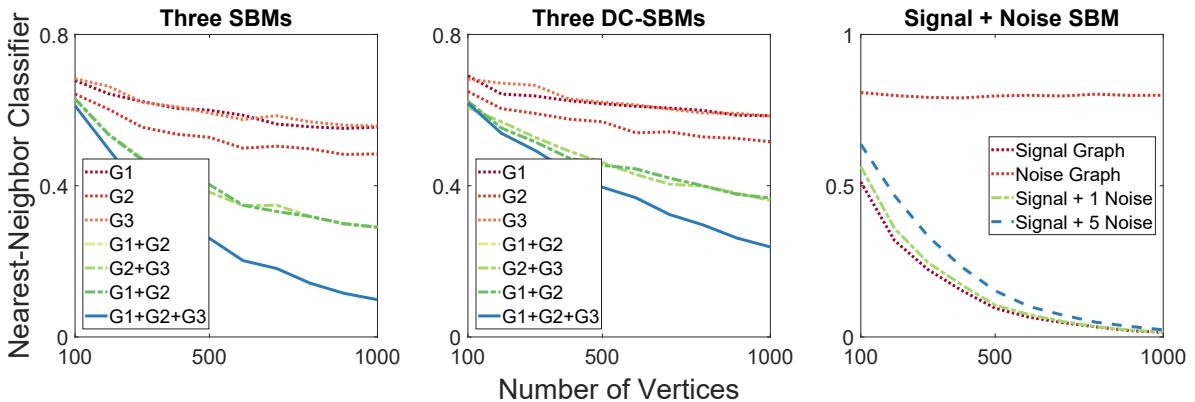


Figure 3: The figure shows the 5-fold classification error of graph fusion embedding, averaged over 20 Monte-Carlo replicates, for the three simulation settings.

6. Real Data

In this section, we showcase the performance and synergistic effect of the graph fusion embedding using a wide range of real-world graphs. The dataset consists of various types of graphs, including binary graphs, weighted graphs, directed and undirected graphs, as well as general graphs such as distance matrices of vertex attributes. The majority of the data were sourced from the network repository [26], while a few were collected directly by the authors. All the public data has been properly formatted in MATLAB and is made

available on GitHub². The results for Omnibus, MASE, and USE methods are provided in the appendix for comparison.

6.1. Two-Graph Data

Table 1 presents the 5-fold classification error for five graph datasets, each having two graphs. The results are reported as the mean and standard deviation over 20 Monte-Carlo replicates. A brief description of each data and pre-processing steps is as follows:

- The C-elegans neuron data [20] consists of two binary graphs with 253 vertices, over 1000 edges, and a label vector with three classes.
- The cora dataset [17] is a citation network with 2708 vertices and 5429 edges and 7 classes, and each vertex has a 1433-dimensional 0/1 vertex attribute, indicating the absence or presence of corresponding words from the dictionary. \mathbf{A}_1 is a binary matrix indicating the existence of edges, while $\mathbf{A}_2 = \mathbf{1} - \mathbf{D}$ where \mathbf{D} is the pairwise cosine distance matrix of the vertex attributes and $\mathbf{1}$ is the matrix of ones.
- The citeseer data [8] is another citation network with 3312 vertices and 4715 edges and 6 classes, with a vertex attribute of 3703 dimensions. We apply the same processing as in the cora dataset to yield a binary graph \mathbf{A}_1 and a general graph $\mathbf{A}_2 = \mathbf{1} - \mathbf{D}$ from the cosine distance of the attributes.
- The IMDB [41] data consists of two slightly un-matched graphs, one with binary labels and 19773 vertices, and another with three-class labels and 19502 vertices. We remove the 271 unmatched vertices to ensure a common vertex set, and use the 3-class label as the label vector.
- COIL-RAG has 11757 vertices with 23588 edges, 100 classes, and a vertex attribute of 64 dimensions. \mathbf{A}_1 is a binary matrix indicating the existence of edges, while $\mathbf{A}_2 = \max\{\mathbf{D}\} - \mathbf{D}$ where \mathbf{D} is the pairwise Euclidean distance matrix of the vertex attributes.

In the cora and citeseer datasets, we employed the cosine distance as it is a widely-used and effective distance measure in word-frequency analysis [4]. Additionally, we transformed the distance matrix by subtracting the distance entries from the maximum value in the distance matrix, denoted as $\max\{\mathbf{D}\}$ (which equals 1 for cosine distance). This transformation ensures that the resulting matrix is a valid kernel or similarity matrix [29].

The results presented in Table 1 (and compared with appendix Table B.2) clearly demonstrates that the graph fusion embedding outperforms all other methods by a significant margin across all real datasets, regardless of whether it involves single graph classification, multiple binary graph classification, or multiple

²<https://github.com/cshen6/GraphEmd>

general graphs. The only exception is the C-elegans data, where the fusion embedding exhibits slightly lower performance, likely due to the small size of the vertex set. Furthermore, the graph fusion embedding consistently exhibits improved classification performance as more graphs are included.

	Graph 1	Graph 2	Graph 1+2
C-elegans	$42.1\% \pm 2.3\%$	$44.5\% \pm 3.0\%$	$38.3\% \pm 2.4\%$
Cora	$19.2\% \pm 0.5\%$	$29.7\% \pm 0.5\%$	$15.8\% \pm 0.5\%$
Citeseer	$32.3\% \pm 1.6\%$	$31.1\% \pm 0.3\%$	$26.1\% \pm 0.5\%$
COIL-RAG	$7.44\% \pm 0.4\%$	$48.4\% \pm 0.03\%$	$3.72\% \pm 0.2\%$
IMDB	$0.9\% \pm 2\%$	$0.02\% \pm 0.02\%$	$0\% \pm 0\%$

Table 1: Evaluate the classification errors using graph fusion embedding. We report the average 5-fold errors and the standard deviation over 20 Monte-Carlo replicates.

6.2. Three-Graph and Three-Label Data

The letter dataset consists of three binary graphs, named letter-high, letter-med, and letter-low, each of which has its own label vector with a total of 15 classes. The graphs contain over 10500 vertices and 20250, 14092 and 14426 edges, respectively. After removing several unmatched vertices, we use all three graphs to predict the class labels, and the results are shown in Table 2. The results show that the fusion embedding consistently exhibits the synergistic effect. Specifically, the classification error for each graph and its corresponding label is better than that of a mismatched pair, for example using Letter-High to predict Label-High results in better error than using Letter-Mid for the same task. Additionally, using all three graphs yields the best classification error in all cases.

The dataset comprises graph representations of 15 letters: a, e, f, h, i, k, l, m, n, t, v, w, x, y, z. The labels "high," "med," and "low" denote the letter's shape, with "high" indicating capitalized letters. The drawings are distorted, making letter recognition a non-trivial task. In the graph, vertices denote the endpoints of the letter drawing, while edges represent the lines in the drawing. Our results indicate that capitalized letters are the easiest to recognize, likely due to their distinct shapes compared to smaller letters. Moreover, employing the fusion embedding using all three graphs yields the best performance in identifying the letter (or, more precisely, the endpoint of a letter) each vertex belongs to.

6.3. Four-Graph Wikipedia Data

The Wikipedia dataset consists of four disparate graph representations we collected from Wikipedia articles [30, 31]. The dataset has a total of $n = 1382$ Wikipedia articles based on a 2-neighborhood of the

	Letter-High	Letter-Med	Letter-Low	All Graphs
Label-High	7.86% \pm 0.3%	44.2% \pm 0.8%	43.0% \pm 0.9%	6.09% \pm 0.2%
Label-Med	44.6% \pm 0.9%	14.4% \pm 0.4%	44.1% \pm 0.6%	9.75% \pm 0.3%
Label-Low	44.3% \pm 0.6%	43.2% \pm 0.5%	13.9% \pm 0.4%	9.50% \pm 0.2%

Table 2: Evaluate the classification errors using fusion embedding for a three-graph three-label data.

English article “Algebraic Geometry” and the corresponding French articles. Latent Dirichlet Allocation is applied to the word-frequency of each article to calculate the English text feature and French text feature, from which we generate two graphs \mathbf{A}_1 and \mathbf{A}_2 via the cosine distance transformation. The hyperlink information in the English articles and the corresponding French articles yields \mathbf{A}_3 and \mathbf{A}_4 . All articles are manually labeled into $K = 5$ disjoint classes: category, people, locations, date, math.

Table 3 displays the classification errors using 5-nearest-neighbor on graph fusion embedding for the Wikipedia dataset. For each individual graph, it is evident that the text graphs provide more information for predicting the class of the article, as expected. On the other hand, the hyperlink graph exhibits significantly higher classification errors, indicating less signals. Moreover, combining the text and hyperlink graphs results in slightly improved performance compared to using individual graph. The most favorable outcome is achieved when all four graphs are utilized, showcasing the synergistic effect.

	English Text	French Text	English Hyperlinks	French Hyperlinks
One Graph	19.3% \pm 0.5%	18.5% \pm 0.5%	48.6% \pm 0.9%	52.9% \pm 1.2%
Two Graphs	16.2% \pm 0.5% (1+2)	18.4% \pm 0.6% (2+4)	19.0% \pm 0.5% (1+3)	44.7% \pm 1.0% (3+4)
Three Graphs	16.2% \pm 0.5% (-4)	16.1% \pm 0.5% (-3)	18.8% \pm 0.5% (-2)	18.4% \pm 0.5% (-1)
All			16.0% \pm 0.5%	

Table 3: Evaluate the classification errors using graph fusion embedding for the Wikipedia data. The first row evaluates the embedding of each individual graph; the second row evaluates the embedding of using two graphs, e.g., 1+2 means English text + French text; the third row evaluates the embedding of using three graphs, e.g., -4 means all graphs excluding French hyperlinks; the last row uses all four graphs.

7. Discussion

This paper presented the graph fusion embedding, and provided compelling evidence on its theoretical soundness and numerical advantages. The most important feature of this method is the synergistic effect, which results in improved numerical performance as more graphs are included. To some extent, the method

can be viewed as a supervised and deterministic version of spectral embedding, but without requiring any additional transformation, and works well for general graphs. In contrast, many competitor methods require extra alignment or additional singular value decomposition.

This paper suggests several promising directions for future research. Firstly, as the dimension of the fusion embedding always equals MK , for an extremely large number of graphs with high M , or an extremely large number of classes M , exploring dimension reduction techniques and strategies for selecting graphs would be intriguing. Secondly, since the graph encoder embedding can be employed for distance and kernel matrices, exploring connections with other areas such as multi-dimensional scaling [6, 7], distance and kernel correlation [37, 10, 32], etc., could provide valuable insights. Lastly, the application of this method to general graphs suggests its potential to enhance existing data fusion tasks, provided an appropriate metric can be identified. Hence, exploring more applications using this method and possible variants on different data regimes would be an exciting avenue for future research.

Acknowledgment

This work was supported in part by the Defense Advanced Research Projects Agency under the D3M program administered through contract FA8750-17-2-0112, the National Science Foundation HDR TRIPODS 1934979, the National Science Foundation DMS-2113099, and by funding from Microsoft Research.

References

- [1] L. Akoglu, H. Tong, and D. Koutra. Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2015.
- [2] J. Arroyo, A. Athreya, J. Cape, G. Chen, C. E. Priebe, and J. T. Vogelstein. Inference for multiple heterogeneous networks with a common invariant subspace. *Journal of Machine Learning Research*, 22(142):1–49, 2021.
- [3] A.-L. Barabási and Z. N. Oltvai. Network biology: Understanding the cell’s functional organization. *Nature Reviews Genetics*, 5(2):101–113, 2004.
- [4] D. Blei, A. Ng, and M. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [5] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4-5):175–308, 2006.
- [6] I. Borg and P. Groenen. *Modern Multidimensional Scaling: theory and applications*. Springer-Verlag, 2005.

- [7] T. Cox and M. Cox. *Multidimensional Scaling*. Chapman and Hall, 2001.
- [8] C. L. Giles, K. D. Bollacker, and S. Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the Third ACM Conference on Digital Libraries*, pages 89–98, 1998.
- [9] M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proceedings of National Academy of Science*, 99(12):7821–7826, 2002.
- [10] A. Gretton, R. Herbrich, A. Smola, O. Bousquet, and B. Scholkopf. Kernel methods for measuring independence. *Journal of Machine Learning Research*, 6:2075–2129, 2005.
- [11] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.
- [12] B. Karrer and M. E. J. Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83:016107, 2011.
- [13] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [14] R. Liu and A. Krishnan. Pecanpy: a fast, efficient and parallelized python implementation of node2vec. *Bioinformatics*, 37(19):3377–3379, 2021.
- [15] V. Lyzinski, D. Fishkind, and C. E. Priebe. Seeded graph matching for correlated Erdos-Renyi graphs. *Journal of Machine Learning Research*, 15:3513–3540, 2014.
- [16] V. Lyzinski, Y. Park, C. E. Priebe, and M. W. Trosset. Fast embedding for jofc using the raw stress criterion. *Journal of Computational and Graphical Statistics*, 26(4):786–802, 2017.
- [17] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3:127–163, 2000.
- [18] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.
- [19] Y. Park, C. E. Priebe, and A. Youssef. Anomaly detection in time series of graphs using fusion of graph invariants. *IEEE Journal of Selected Topics in Signal Processing*, 7:67–75, 2013.
- [20] D. Pavlovic, P. Vertes, E. Bullmore, W. Schafer, and T. Nicholas. Stochastic blockmodeling of the modules and core of the *caenorhabditis elegans* connectome. *PLoS ONE*, 9(9):e97584, 2014.

[21] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710. ACM, 2014.

[22] C. E. Priebe, D. J. Marchette, Z. Ma, and S. Adali. Manifold matching: Joint optimization of fidelity and commensurability. *Brazilian Journal of Probability and Statistics*, 27(3):377–400, 2013.

[23] W. M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971.

[24] S. Ranshous, S. Shen, D. Koutra, S. Harenberg, C. Faloutsos, and N. F. Samatova. Anomaly detection in dynamic networks: a survey. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(3):223–247, 2015.

[25] K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *Annals of Statistics*, 39(4):1878–1915, 2011.

[26] R. A. Rossi and N. K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. URL <https://networkrepository.com>.

[27] P. Rubin-Delanchy, J. Cape, M. Tang, and C. E. Priebe. A statistical interpretation of spectral embedding: The generalised random dot product graph. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(4):1446–1473, 2022.

[28] C. Shen. Encoder embedding for general graph and node classification. *arXiv preprint arXiv:2405.15473*, 2024.

[29] C. Shen and J. T. Vogelstein. The exact equivalence of distance and kernel methods in hypothesis testing. *AStA Advances in Statistical Analysis*, 105(3):385–403, 2021.

[30] C. Shen, M. Sun, M. Tang, and C. E. Priebe. Generalized canonical correlation analysis for classification. *Journal of Multivariate Analysis*, 130:310–322, 2014.

[31] C. Shen, J. T. Vogelstein, and C. Priebe. Manifold matching using shortest-path distance and joint neighborhood selection. *Pattern Recognition Letters*, 92:41–48, 2017.

[32] C. Shen, S. Panda, and J. T. Vogelstein. The chi-square test of distance correlation. *Journal of Computational and Graphical Statistics*, 31(1):254–262, 2022.

[33] C. Shen, Y. Park, and C. E. Priebe. Graph encoder ensemble for simultaneous vertex embedding and community detection. In *2023 2nd International Conference on Algorithms, Data Mining, and Information Technology*. ACM, 2023.

- [34] C. Shen, Q. Wang, and C. E. Priebe. One-hot graph encoder embedding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7933 – 7938, 2023.
- [35] C. Shen, J. Larson, H. Trinh, X. Qin, Y. Park, and C. E. Priebe. Discovering communication pattern shifts in large-scale labeled networks using encoder embedding and vertex dynamics. *IEEE Transactions on Network Science and Engineering*, 11(2):2100 – 2109, 2024.
- [36] D. Sussman, M. Tang, D. Fishkind, and C. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107(499):1119–1128, 2012.
- [37] G. Szekely and M. Rizzo. Brownian distance covariance. *Annals of Applied Statistics*, 3(4):1233–1303, 2009.
- [38] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*, 2011.
- [39] L. Varshney, B. Chen, E. Paniagua, D. Hall, and D. Chklovskii. Structural properties of the caenorhabditis elegans neuronal network. *PLoS Computational Biology*, 7(2):e1001066, 2011.
- [40] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32:4–24, 2019.
- [41] P. Yanardag and S. Vishwanathan. Deep graph kernels. In *KDD 15: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1365–1374, 2015.
- [42] S. Young and E. Scheinerman. Random dot product graph models for social networks. In *Algorithms and Models for the Web-Graph*, pages 138–149. Springer Berlin Heidelberg, 2007.
- [43] Y. Zhao, E. Levina, and J. Zhu. Consistency of community detection in networks under degree-corrected stochastic block models. *Annals of Statistics*, 40(4):2266–2292, 2012.

APPENDIX

Appendix A. Theorem Proofs

Theorem 1. Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. Let n be the number of vertices with known labels. Then, for any vertex i belonging to class \mathbf{Y}_i , its graph fusion embedding satisfies:

$$\|\tilde{\mathbf{Z}}(i, \cdot) - \tilde{\mathbf{B}}(\mathbf{Y}_i, \cdot)\| = O\left(\frac{1}{\sqrt{n}}\right) \xrightarrow{n \rightarrow \infty} 0.$$

Proof. First, based on the convergence theorem in [34], under the DC-SBM model, the un-normalized embedding for vertex i of class \mathbf{Y}_i satisfies

$$\mathbf{Z}_m(i, k) \xrightarrow{n \rightarrow \infty} \theta_i E(F_\theta) \mathbf{B}_m(\mathbf{Y}_i, k).$$

It is important to note that the above conditioning assumes knowledge of everything about vertex i , including the degree θ_i and class label \mathbf{Y}_i . It follows that the Euclidean norm satisfies

$$\begin{aligned} \|\mathbf{Z}_m(i, \cdot)\|^2 &= \left(\sum_{k=1}^K \mathbf{Z}_m^2(i, k)\right) \\ &\xrightarrow{n \rightarrow \infty} \theta_i^2 E^2(\theta_i) \left(\sum_{k=1}^K \mathbf{B}_m^2(\mathbf{Y}_i, k)\right) \\ &= \theta_i^2 E^2(\theta_i) \|\mathbf{B}_m(\mathbf{Y}_i, \cdot)\|^2 \end{aligned}$$

Then, the normalized embedding satisfies

$$\tilde{\mathbf{Z}}_m(i, k) = \frac{\mathbf{Z}_m(i, k)}{\|\mathbf{Z}_m(i, \cdot)\|} \xrightarrow{n \rightarrow \infty} \frac{\mathbf{B}_m(\mathbf{Y}_i, k)}{\|\mathbf{B}_m(\mathbf{Y}_i, \cdot)\|} = \tilde{\mathbf{B}}_m(\mathbf{Y}_i, k).$$

By concatenating every dimension from 1 to K within graph m , we obtain:

$$\|\tilde{\mathbf{Z}}_m(i, \cdot) - \tilde{\mathbf{B}}_m(\mathbf{Y}_i, \cdot)\| \xrightarrow{n \rightarrow \infty} 0.$$

The graph fusion embedding is a row-concatenation of $\tilde{\mathbf{Z}}_m(i, \cdot)$ for $m = 1, \dots, M$, so it follows that

$$\|\tilde{\mathbf{Z}}(i, \cdot) - \tilde{\mathbf{B}}(\mathbf{Y}_i, \cdot)\| \xrightarrow{n \rightarrow \infty} 0$$

where $\tilde{\mathbf{B}}$ is the row concatenation of every $\tilde{\mathbf{B}}_m$.

To verify the rate of convergence, note that

$$\mathbf{W}(i, k) = 1/n_k$$

when $\mathbf{Y}_i = k$, so

$$\mathbf{Z}_m(i, k) = \mathbf{A}_m(i, 1 : n) \mathbf{W}(1 : n, k) = \sum_{j=1, \dots, n}^{\mathbf{Y}_j=k} \frac{\mathbf{A}_m(i, j)}{n_k}.$$

Therefore, each dimension is essentially a sample mean using n_k observations per class, with a standard deviation of $O(\frac{1}{\sqrt{n_k}})$. Hence, the convergence rate for dimension k is $O(\frac{1}{\sqrt{n_k}})$, and it is the same rate for any m . Now, if we consider the Euclidean norm of all dimensions, i.e.,

$$\|\mathbf{Z}_m(i, :)\| = \sqrt{\left(\sum_{k=1}^K \mathbf{Z}_m^2(i, k)\right)},$$

its convergence rate is determined by the slowest across all dimensions, which ends up being $O(\max_{k=1, \dots, K} \frac{1}{\sqrt{n_k}})$.

Note that the probabilistic model assumed that all classes are non-trivial and $\pi_k > 0$ for all k , it follows that $n_k = O(n)$ for all k , and the overall convergence rate can be simplified to $O(\frac{1}{\sqrt{n}})$. \square

Theorem 2. *Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. As the number of training vertices increases to infinity, the graph fusion embedding achieves asymptotically perfect vertex classification, i.e.,*

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) = 0,$$

if and only if there are no repeating rows in $\tilde{\mathbf{B}}$.

Proof. First, it is important to note that the convergence described in Theorem 1 applies to both training vertices and testing vertices, as both use the same W matrix that is based on the labels of the training vertices only. Consequently, the embedding for a testing vertex of class k asymptotically equals the embedding for all training vertices of class k , which in turn converges to $\tilde{\mathbf{B}}(k, :)$.

Perfect classification is achievable only when the vertex embedding of each class are completely distinct and well-separated from the vertex embedding of all other classes. Consequently, the testing vertices can be classified perfectly if and only if $\tilde{\mathbf{B}}(k, :) \neq \tilde{\mathbf{B}}(l, :)$ for all pairs of classes $k \neq l$. This condition is met if and only if every row of $\tilde{\mathbf{B}}$ is unique. \square

Theorem 3. *Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the DC-SBM model in Section 3.1. For any subset of graphs $M_1 \leq M$, the classification error using the fusion embedding satisfies:*

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) - L_n(\mathbf{A}_1, \dots, \mathbf{A}_{M_1}) \leq 0$$

for sufficiently large sample size n .

Proof. As the graph fusion embedding converges to the block probability matrix, it suffices to consider the concatenated block matrices. Let $\tilde{\mathbf{B}}^{M_1} \in \mathbb{R}^{K \times M_1 K}$ denote the concatenated block matrix for M_1 graphs, and $\tilde{\mathbf{B}}^{M-M_1} \in \mathbb{R}^{K \times (M-M_1)K}$ as the concatenated block matrix for the remaining $M - M_1$ graphs, i.e.,

$$\tilde{\mathbf{B}}^M = [\tilde{\mathbf{B}}^{M_1}, \tilde{\mathbf{B}}^{M-M_1}].$$

The graph fusion embedding using only M_1 graphs asymptotically equals $\tilde{\mathbf{B}}^{M_1}$, and the embedding using M graphs asymptotically equals $\tilde{\mathbf{B}}^M$.

If there are no repeating rows in $\tilde{\mathbf{B}}^{M_1}$, then the classification is already perfect using the fusion embedding of M_1 graphs. Moreover, all K rows will remain unique in $\tilde{\mathbf{B}}^M$, ensuring that the classification remains perfect when using the fusion embedding of all M graphs. Therefore, for sufficiently large n , using more graphs in the fusion embedding cannot deteriorate the classification performance.

On the other hand, if there exists a pair $k \neq l$ such that

$$\tilde{\mathbf{B}}^{M_1}(k, :) = \tilde{\mathbf{B}}^{M_1}(l, :),$$

i.e., using the fusion embedding of M_1 graphs, vertices in class k cannot be separated from vertices in class l . However, if the remaining $M - M_1$ graphs can separate these two classes, we have

$$\tilde{\mathbf{B}}^{M-M_1}(k, :) \neq \tilde{\mathbf{B}}^{M-M_1}(l, :)$$

and

$$\tilde{\mathbf{B}}^M(k, :) \neq \tilde{\mathbf{B}}^M(l, :).$$

Namely, these two classes may become perfectly separable using the fusion embedding of all M graphs. Therefore, for sufficiently large n , adding additional graphs never deteriorates the classification performance, but has the potential to improve and successfully classify previously mis-classified vertices. Consequently,

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_{M_1}, \dots, \mathbf{A}_M) - L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) \leq 0.$$

□

Theorem 4. Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the general graph model in Section 4.1. Let n be the number of vertices with known labels. Then, for any vertex i belonging to class \mathbf{Y}_i , its graph fusion embedding satisfies:

$$\|\tilde{\mathbf{Z}}(i, :) - \mu(i)\| = O\left(\frac{1}{\sqrt{n}}\right) \xrightarrow{n \rightarrow \infty} 0.$$

Proof. For each $\mathbf{Z}_m(i, k)$:

$$\mathbf{Z}_m(i, k) = \mathbf{A}_m(i, 1 : n) \mathbf{W}(1 : n, k) = \sum_{j=1, \dots, n}^{\mathbf{Y}_j=k} \frac{\mathbf{A}_m(i, j)}{n_k}.$$

Given vertex i , each $\mathbf{A}_m(i, 1), \mathbf{A}_m(i, 2), \dots, \mathbf{A}_m(i, n)$ are independently and identically distributed, and they have finite moments. As n increases to infinity, so is n_k , and law of large numbers applies:

$$\mathbf{Z}_m(i, k) \xrightarrow{n \rightarrow \infty} E(\mathbf{A}_m(i, j) | i, \mathbf{Y}_j = k).$$

Then the Euclidean norm satisfies

$$\|\mathbf{Z}_m(i, :)\| \xrightarrow{n \rightarrow \infty} \|E(\mathbf{A}_m(i, j) | i)\|,$$

and

$$\tilde{\mathbf{Z}}_m(i, k) = \frac{\mathbf{Z}_m(i, k)}{\|\mathbf{Z}_m(i, :)\|} \xrightarrow{n \rightarrow \infty} \frac{E(\mathbf{A}_m(i, j) | i, \mathbf{Y}_j = k)}{\|E(\mathbf{A}_m(i, j) | i)\|}.$$

By concatenating all dimensions within the m th graph, we obtain:

$$\|\tilde{\mathbf{Z}}_m(i, :) - \mu_m(i)\| \xrightarrow{n \rightarrow \infty} 0.$$

Finally, by concatenating all individual graph embedding into the fusion embedding, we have:

$$\|\tilde{\mathbf{Z}}(i, :) - \mu(i)\| \xrightarrow{n \rightarrow \infty} 0.$$

Note that the convergence rate can be proved in the same manner as in the proof of Theorem 1, and thus is not repeated here. \square

Theorem 5. *Suppose that $\{\mathbf{A}_m, m = 1, \dots, M\}$ follows the general graph model in Section 4.1. As the number of training vertices increases to infinity, the graph fusion embedding achieves asymptotically perfect vertex classification, i.e.,*

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) = 0$$

if and only if

$$\sum_{1 \leq k < l \leq K} h_k(x)h_l(x) = 0 \tag{A.1}$$

for any $x \in \mathbb{R}^{MK}$.

Additionally, for any subset of graphs $M_1 \leq M$, the classification error using the fusion embedding satisfies:

$$\lim_{n \rightarrow \infty} L_n(\mathbf{A}_1, \dots, \mathbf{A}_{M_1}, \dots, \mathbf{A}_M) - L_n(\mathbf{A}_1, \dots, \mathbf{A}_M) \leq 0$$

for sufficiently large sample size n .

Proof. As the fusion embedding converges to the concatenated class-conditional means, it suffices to focus on the density family $\{h_k(x), k = 1, \dots, K\}$ to analyze the classification performance. When all densities are distinct from each other, it implies that vertex embedding for each class are asymptotically separated from embedding of other classes, leading to perfect classification. In other words, the embedding in class k is perfectly distinguishable from the embedding in class l if and only if $h_k(x)h_l(x) = 0$. This condition extends to Equation A.1 by considering all possible pairs.

Regarding the synergistic property, let $\{h_k^{M_1}, k = 1, \dots, K\}$ represents the densities using M_1 graphs, and $\{h_k^M, k = 1, \dots, K\}$ represents the densities using all M graphs. Assuming Equation A.1 is not equal to zero for embedding using M_1 graphs, we define:

$$\tau^{M_1} = \{x \in \mathbb{R}^{M_1 K}, \sum_{1 \leq k < l \leq K} h_k^{M_1}(x)h_l^{M_1}(x) > 0\}$$

as the set of mis-classified points. An element $u \in \tau^{M_1}$ has at least one pair of (k, l) such that $h_k^{M_1}(u)h_l^{M_1}(u) > 0$, meaning that u may occur in both class k and class l , making these two classes not perfectly separable.

Now, upon adding an additional $M - M_1$ graphs, the element u is appended more entries from the additional graphs, leading to $[u, u'] \in \mathbb{R}^{MK}$. Either $[u, u'] \in \tau^M$ and it remains a mis-classified point, or u' is perfectly separable based on the additional $M - M_1$ graphs, leading to $h_k^M([u, u'])h_l^M([u, u']) = 0$ for the fusion embedding using M graphs. On the other hand, any element not in τ^{M_1} is already separable in the $\mathbb{R}^{M_1 K}$, and still separable in the larger, appended \mathbb{R}^{MK} space. Therefore, we have:

$$\tau^M \subseteq \tau^{M_1},$$

indicating that the probability of mis-classification is either the same or reduced as more graphs are included in the fusion embedding. \square

Appendix B. Additional Experiments

Appendix B.1. Heterogeneous Graphs

Figure B1 illustrates the synergistic effect on the same simulations as in Figure 3, but with heterogeneous graphs. In the SBM setting, we used the same three-graph models as in simulation 1, but we introduced heterogeneity by multiplying each adjacency matrix A_m by a random factor c , where $c \sim \text{Uniform}(0, 0.5)$ for $m = 1, 2, 3$. Additionally, we included 17 noise graphs that are independent of the labels, with $A_m^h = A_m * c$, where $c \sim \text{Uniform}(1, 2)$ for $m > 3$. Consequently, the degrees of the vertices are randomized and no longer the same across different graphs, and the noise graphs have larger degrees than the signal graphs. Despite this heterogeneity, Figure B1 clearly shows the same synergistic effect: using all three graphs yields better performance than using any single graph, and the inclusion of additional noise graphs only minimally deteriorates the error. This observation holds true for the DC-SBM setting as well. Note that we omitted the two-graph fusion embedding in the figure for clearer visualization.

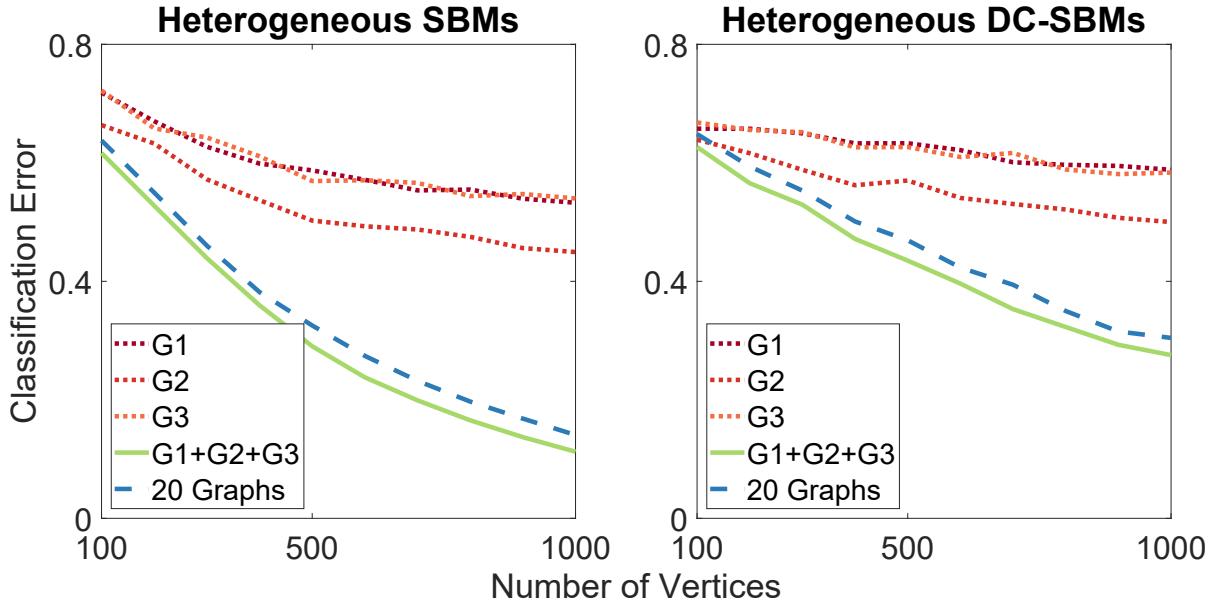


Figure B1: The same experiments as in Figure 3 were conducted using three heterogeneous signal graphs, and the blue dotted line shows the error for adding 17 noise graphs to the three signal graphs.

Appendix B.2. Alternative Classifiers

To show that the synergistic effect occurs regardless of the classifier being used, Figure B2 shows the classification results on the fusion embedding using three different classifiers: a two-layer neural network with a neuron size of 100, the 5-nearest-neighbor classifier (as used in the main paper), and linear discriminant analysis. While there are slight variations in the actual classification errors among the classifiers, the synergistic effect remains consistent across all of them, i.e., including more graphs consistently improves the classification error, and the presence of noise graphs does not deteriorate the error for large sample sizes.

Appendix B.3. Clustering

In this subsection, we introduce an experimental clustering algorithm for multiple graphs in the absence of ground-truth labels. The approach parallels the clustering algorithm designed for a single graph [33].

Since ground-truth labels are unavailable, we begin by randomly initializing a label vector. This randomized label vector, along with the multiple graphs, serves as input to the main method outlined in Section 2. Subsequently, we apply K-means clustering on the output embedding to obtain a new label vector. This new label vector, along with the multiple graphs, is then fed back into the main algorithm iteratively. This iterative process continues until either the new and previous labels are identical, or a pre-defined iteration limit is reached (set to 20 in our experiments).

The clustering performance is evaluated using the adjusted rand index (ARI), a popular metric for measuring the similarity between two label vectors of the same size. The ARI ranges from $-\infty$ to 1, where

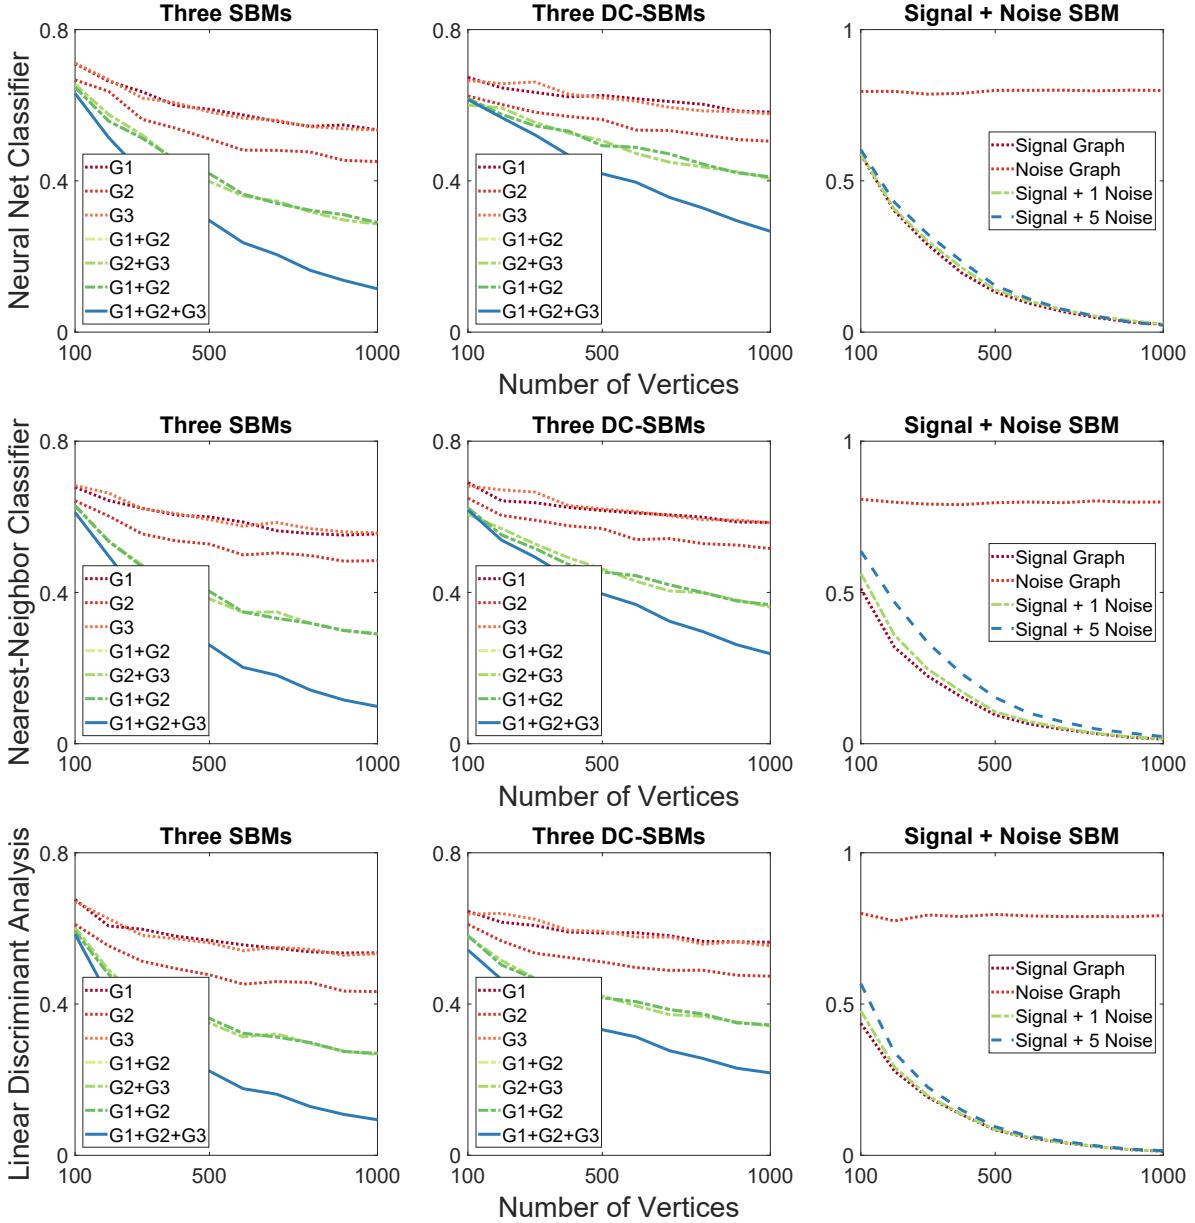


Figure B2: The same experiments as in Figure 3 were conducted using three different classifiers: neural network (first row), nearest-neighbor (second row), and linear discriminant (third row).

a higher positive value indicates better match quality and a score of 1 denotes a perfect performance [23]. Table B.1 presents the clustering performance on individual graphs, all three signal graphs, and all three signal graphs plus 17 noise graphs, based on the heterogeneous DC-SBM models in Figure B1.

It is evident from the results that incorporating more graphs leads to improvements in clustering performance. Using all three graphs yields a higher ARI, while the inclusion of noise graphs only minimally affects the performance. Although the synergistic effect appears to extend to the clustering task, there are several areas warranting further investigation. Firstly, due to the iterative nature and random label initialization of the clustering version, a closed-form analysis of embedding convergence is not feasible in the same manner. Secondly, as unsupervised learning poses greater challenges compared to supervised learning, achieving comparable clustering performance requires larger sample sizes. Lastly, as randomization is involved, the algorithm is inherently more volatile and susceptible to noise, as evidenced by the high standard deviation. Therefore, while the clustering performance is promising, further research and algorithm refinement are necessary in the clustering scenario.

	G1	G2	G3	G1+2+3	G1+2+3+Noise
$n = 1000$	0.03 ± 0.02	0.11 ± 0.05	0.02 ± 0.02	0.04 ± 0.02	0.04 ± 0.02
$n = 2000$	0.15 ± 0.03	0.30 ± 0.04	0.14 ± 0.03	0.40 ± 0.10	0.40 ± 0.10
$n = 3000$	0.21 ± 0.02	0.36 ± 0.03	0.21 ± 0.02	0.67 ± 0.17	0.68 ± 0.17
$n = 4000$	0.23 ± 0.02	0.39 ± 0.03	0.23 ± 0.01	0.88 ± 0.10	0.85 ± 0.13

Table B.1: The table displays the adjusted rand index (ARI) for the clustering version of the graph fusion embedding, based on the simulations in the right panel of Figure B1. We conducted 20 Monte Carlo replicates and report the average ARI along with the standard deviations.

Appendix B.4. Additional Benchmarks

In this section, we reproduce the simulations and experiments from the main paper for the Omnibus, MASE, and USE methods. Since these methods rely on spectral embedding, they require selecting a dimension parameter d . We evaluate all possible values of d from 1 to 30 and report the best error achieved. The classifier used throughout is the 5-nearest-neighbor classifier, consistent with the main paper.

The simulation Figure B3 demonstrate that all three methods are capable of improving classification errors when the additional graph contains signal for the labels. Among the three methods, USE performs the best and has a similar error as the fusion embedding in the main paper. However, in simulation 3, all three methods show a significant degradation in error, especially when five independent graphs are included.

For the real data experiments presented in the remaining tables, except for the C-elegans data, all three

methods exhibit significantly worse actual classification error compared to the fusion embedding and do not demonstrate the synergistic effect.

Omnibus	Graph 1	Graph 2	Graph 1+2
C-elegans	$38.7\% \pm 1.3\%$	$43.5\% \pm 1.2\%$	$35.0\% \pm 1.4\%$
Cora	$31.6\% \pm 0.7\%$	$41.2\% \pm 0.9\%$	$30.7\% \pm 0.5\%$
Citeseer	$45.6\% \pm 0.7\%$	$37.5\% \pm 0.5\%$	$33.7\% \pm 0.5\%$
COIL-RAG	$76.2\% \pm 0.9\%$	$36.6\% \pm 0.5\%$	$23.9\% \pm 0.4\%$
IMDB	$19.8\% \pm 0.2\%$	$22.2\% \pm 0.2\%$	$11.8\% \pm 0.2\%$
MASE	Graph 1	Graph 2	Graph 1+2
C-elegans	$39.0\% \pm 1.3\%$	$43.6\% \pm 1.3\%$	$33.6\% \pm 1.3\%$
Cora	$31.3\% \pm 0.4\%$	$41.7\% \pm 0.4\%$	$38.2\% \pm 0.6\%$
Citeseer	$45.8\% \pm 0.4\%$	$37.6\% \pm 0.3\%$	$36.6\% \pm 0.3\%$
COIL-RAG	$77.5\% \pm 0.9\%$	$36.3\% \pm 0.5\%$	$23.7\% \pm 0.4\%$
IMDB	$19.8\% \pm 0.2\%$	$22.2\% \pm 0.2\%$	$45.6\% \pm 0.2\%$
USE	Graph 1	Graph 2	Graph 1+2
C-elegans	$40.1\% \pm 1.5\%$	$44.1\% \pm 1.5\%$	$38.2\% \pm 1.5\%$
Cora	$31.6\% \pm 0.7\%$	$41.2\% \pm 0.9\%$	$30.7\% \pm 0.5\%$
Citeseer	$46.4\% \pm 0.5\%$	$37.3\% \pm 0.3\%$	$36.9\% \pm 0.3\%$
COIL-RAG	$76.9\% \pm 0.9\%$	$35.1\% \pm 0.5\%$	$35.0\% \pm 0.4\%$
IMDB	$19.8\% \pm 0.2\%$	$22.0\% \pm 0.3\%$	$12.6\% \pm 0.1\%$

Table B.2: Same evaluation as in Table 1 using Omnibus, MASE, and USE.

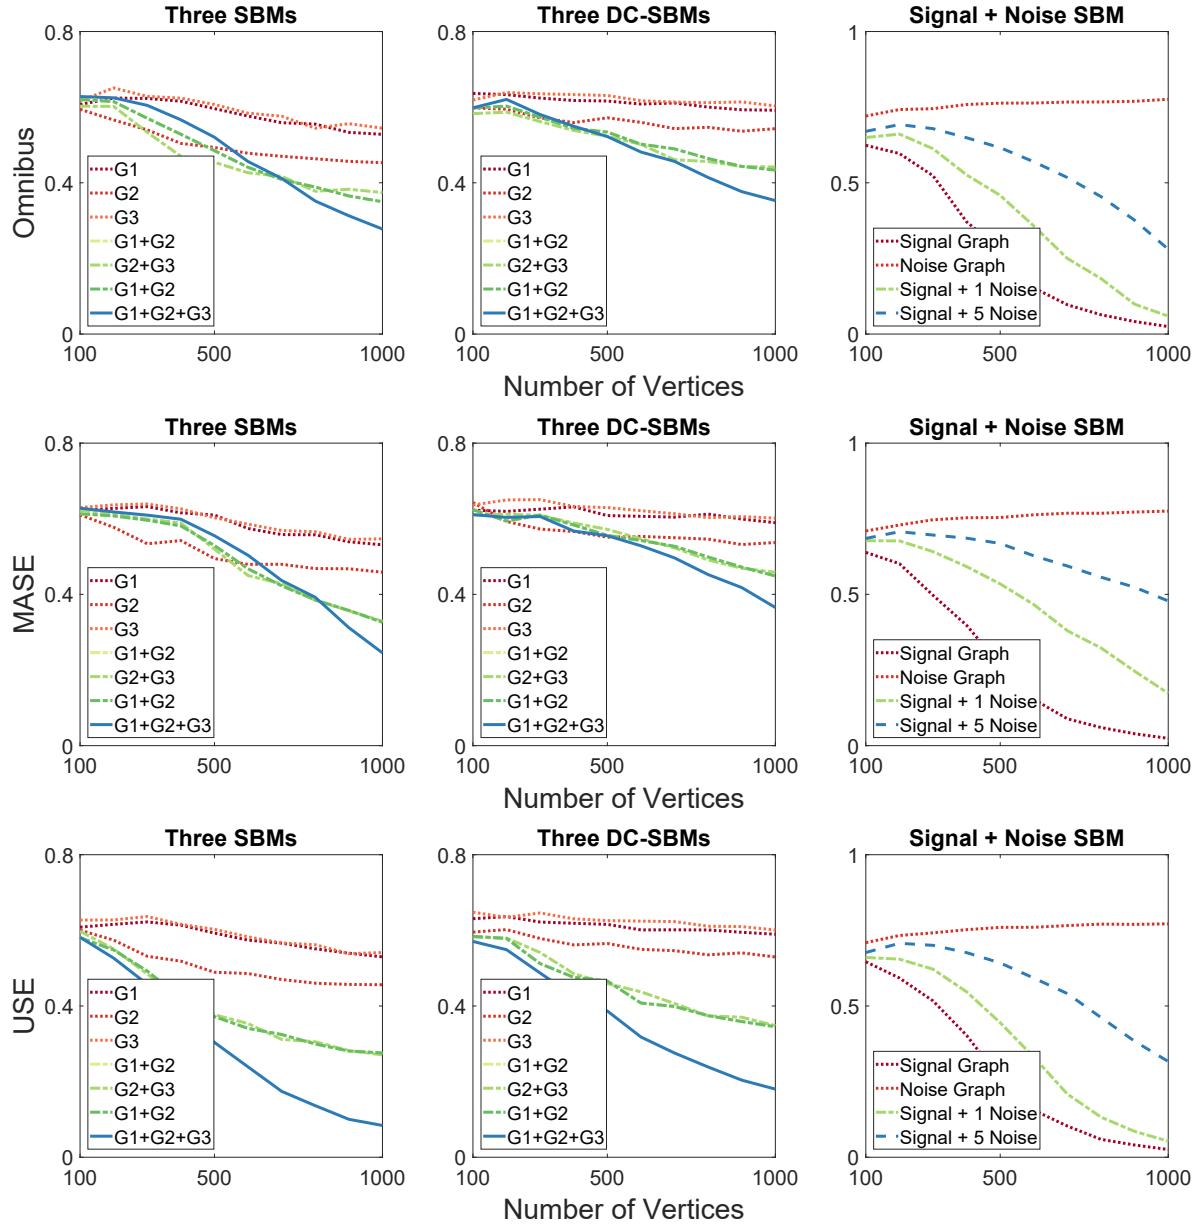


Figure B3: Same experiments as Figure 3 using Omnibus (first row), MASE (second row), and USE (last row).

Omnibus	English Text	English Hyperlinks	French Text	French Hyperlinks
One Graph	$21.3\% \pm 0.5\%$	$22.3\% \pm 0.5\%$	$41.2\% \pm 0.7\%$	$45.5\% \pm 0.6\%$
Two Graphs	$19.4\% \pm 0.4\% (1+2)$	$39.3\% \pm 0.6\% (3+4)$	$41.5\% \pm 0.6\% (1+3)$	$44.4\% \pm 0.5\% (2+4)$
Three Graphs	$34.2\% \pm 0.6\% (-4)$	$37.5\% \pm 0.6\% (-3)$	$39.7\% \pm 0.6\% (-2)$	$39.6\% \pm 0.6\% (-1)$
All	$36.5\% \pm 0.5\%$			
MASE	English Text	English Hyperlinks	French Text	French Hyperlinks
One Graph	$21.8\% \pm 0.4\%$	$23.0\% \pm 0.6\%$	$43.5\% \pm 0.6\%$	$47.2\% \pm 0.5\%$
Two Graphs	$19.0\% \pm 0.5\% (1+2)$	$40.8\% \pm 0.4\% (3+4)$	$41.3\% \pm 0.5\% (1+3)$	$44.2\% \pm 0.7\% (2+4)$
Three Graphs	$38.1\% \pm 0.7\% (-4)$	$39.9\% \pm 0.6\% (-3)$	$39.3\% \pm 0.7\% (-2)$	$39.5\% \pm 0.6\% (-1)$
All	$37.2\% \pm 0.5\%$			
USE	English Text	English Hyperlinks	French Text	French Hyperlinks
One Graph	$21.3\% \pm 0.4\%$	$22.1\% \pm 0.7\%$	$41.4\% \pm 0.7\%$	$45.1\% \pm 0.6\%$
Two Graphs	$19.2\% \pm 0.4\% (1+2)$	$39.9\% \pm 0.6\% (3+4)$	$39.9\% \pm 0.7\% (1+3)$	$43.6\% \pm 0.6\% (2+4)$
Three Graphs	$38.8\% \pm 0.8\% (-4)$	$41.3\% \pm 0.7\% (-3)$	$39.0\% \pm 0.6\% (-2)$	$39.6\% \pm 0.6\% (-1)$
All	$38.7\% \pm 0.6\%$			

Table B.3: Same evaluation as in Table 3 using Omnibus, MASE, and USE.